Towards Internet Interconnection Co-ordination: Measuring Internet Topology

Jan-Pascal van Best
Delft University of Technology
PO Box 5015, 2600 GA Delft, The Netherlands
janpascal@vanbest.org

Abstract

The infrastructure of the Internet consists of many interconnected networks. The free market for Internet connections has led to an enormous growth of the Internet. The maturing of this market may lead to competition issues like those in the telecommunications market. Moreover, there is little understanding on how the interconnection policies of the Internet's constituents (Internet Service Providers or ISPs) determine the topology of the Internet. We have developed a method to create a map of the Internet on the level of ISPs. This map can also be viewed as a map of the interconnection agreements between ISPs, measured from the outside, without asking each and every ISP what other ISPs they interconnect with. When viewed over time, the maps can be used to assess current trends in the Internet infrastructure regarding issues like market power, concentration, dependence upon one or few parties for Internet service, and so on.

Key words: Internet, policy, routing, infrastructure, topology

1. Introduction

The Internet is a complex system. It started out as a research network, fully managed (indirectly) by the U.S. government, but has since evolved into a system of interconnected networks, mostly managed by commercial entities. Vital to the function of the Internet is that all constituent networks are interconnected, either directly or by means of one or more intermediary networks. The focus of the thesis work upon which this paper is based is to unravel the way in which the interconnection policies of the constituent networks determine the resulting topology of the Internet, on the level of IP (Internet Protocol) networks. Insight in this mechanism would contribute to the opportunities of policy makers to determine their approach towards the critical infrastructure the Internet has become for society as a whole. Up until now, policy regarding the Internet infrastructure has been reluctant: the free market in Internet connections has led to an enormous growth of Internet use, and it is feared that interference with the market would result in a slowdown of this process. Working papers of the US Federal Communication Commission (FCC) support the view that even the current trend of concentration in the market or future value-added services will not hinder new entrants to enter the market [Kende00]. Nevertheless, these issues need to be closely watched.

In this paper, we will focus on the Internet's topology and ways to measure that topology. Analysis of the development over time of the results could be used as an early-warning system for concentration leading to monopoly power, thus providing input for policy makers. In the following section, we will discuss the Internet architecture in general and the parameters that are relevant to policy makers regarding issues like market power, concentration, and availability. Next, we will discuss a way to actually measure and analyse the Internet's topology in a meaningful way. After that, we will present the results of these measurements and data analysis and finish with a discussion.

2. Internet Infrastructure

To begin with, we will describe the overall Internet architecture and discuss the constituent components and relations. Then we will discuss what kind of knowledge about the Internet infrastructure is needed for our goals.

2.1. Internet Architecture

Since in this paper we are only interested in the Internet regarded as a large interconnected IP-based network [Kende00], the low-level components of our object of study are *routers* that are connected by internet *links*. Routers and links can be regarded as nodes and edges of a graph. An internet datagram, called a *packet,* travels though the network on its way from its source to its destination. Every router on this trajectory needs to decide to which of its neighbours the packet should be forwarded. This decision process is called *routing*. The packet itself contains no routing information except for its source and destination address. Therefore, each router keeps track of a lookup table that it uses every time it needs to forward a packet, and which tells it to which neighbour a packet must be forwarded based on its destination address. Each Internet router gets the information it needs to fill its lookup table from the router's operator and from other routers. Routing takes place on essentially two levels, based on autonomous systems. An *autonomous system* (AS) is defined as a set of routers that are under a single administrative control and that have a single routing policy. Routing inside an AS, called *interior routing*, is relatively straightforward, since all participating routers are (by definition) under a single administrative control, usually an *Internet Service Provider* (ISP). In the following, we will use the terms AS and ISP interchangeably, where we will mostly use AS when addressing routing issues.

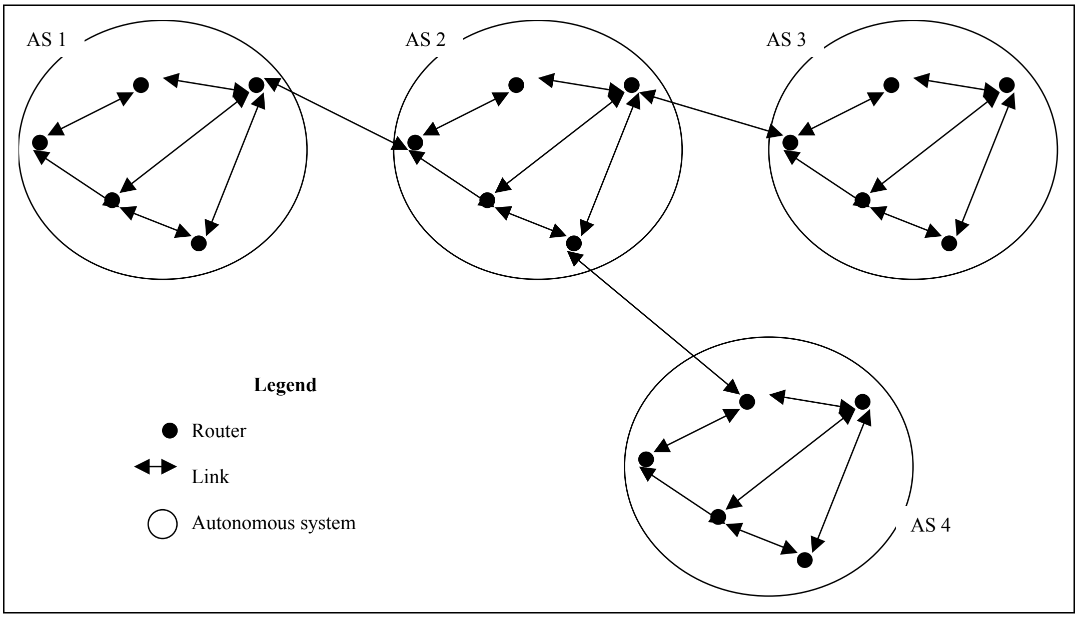

Figure 1: Interconnection graph on router level

Routing that takes place between ASs is called *exterior routing*. For exterior routing, routers on the edges of all autonomous systems, called *border routers*, exchange information about which parts of the Internet can be reached by which routes. This is the point where policy considerations of the individual ISPs step in. ISPs instruct their routers about which IP traffic to accept for which destinations based upon their business policy and their business contracts. This information is propagated over all border routers in the Internet and from them to the interior routers, so that each router in the Internet can determine to which neighbour to forward each packet it is offered. This two-level routing system allows us to view the Internet interconnection graph also in two levels, the IP level (in which nodes are individual routers) and the AS level (in which nodes are autonomous systems). This is illustrated in Figure 1 and Figure 2, which show a fictitious Internet consisting of four ASs, each containing 5 routers. Figure 1 shows the network on the router level. In his example, AS 1, 3 and 4 have a single border router each, while AS 2 has three border routers (connecting AS 2 to the three other ASs). Figure 2 shows the network on the AS level. All routers that belong to the same AS are collapsed into a single node in the graph, giving a graph in which an edge between two ASs means that the ASs exchange Internet traffic.

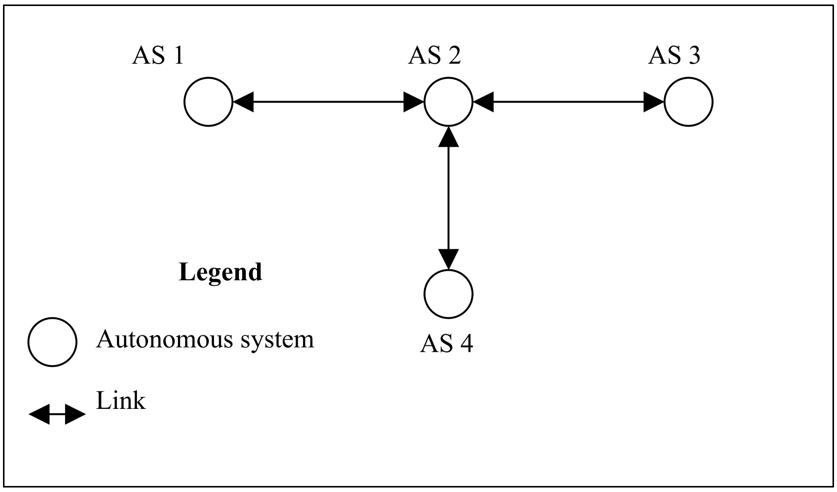

Figure 2: Interconnection graph on AS level

We have seen that ISPs, as autonomous systems, determine what IP traffic they wish to forward for other ISPs by instructing their border routers. ISPs usually do this based on interconnection agreements with other ISPs. The two most common types of interconnection agreements are peering and transit [Norton01]. *Peering* is defined as the business relationship whereby ISPs reciprocally provide to each other connectivity to each other's transit customers; *transit* is the business relationship whereby one ISP provides (usually sells) access to all destinations in its routing table. The set of interconnection agreements that all ISPs maintain determines the topology of the Internet on the AS level. In reverse, we can therefore learn much about the interconnection agreements between ISPs by studying the interconnection graph on the AS level.

2.2. Relevant parameters

Since we are mostly interested in issues of power and concentration in the ISP market, it is natural to look at the number of connections that each ISP has with other ISPs (in mathematical terms, the *degree* of the node representing the AS/ISP). Other useful parameters would be, for each ISP, the number of other ISPs that depend solely on the first ISP for their Internet connection, the size of the ISP in terms of number of routers, or the percentage of shortest paths though the entire network that make use of the ISP's infrastructure. For all these kinds of parameters, an actual map of the Internet on the level of autonomous systems is a first preliminary. The following section will describe how such a map can be obtained.

3. Measuring the Internet

In the previous section, we have argued that a map of the Internet on the level of autonomous systems is necessary in order to determine various parameters of the Internet infrastructure that contribute to our understanding of policy issues like market power and concentration. We will now discuss a number of available data sources and select one of these sources. Then we will describe how the data from this source is analysed in order to produce a map of the Internet on the level of autonomous systems.

3.1. Data source

There are two distinct methods for determining the Internet interconnection graph on the AS level. The first method uses data from border routers to determine which ASs exchange Internet traffic. This method is analysed in [Magoni01]. A major drawback of this method is that border routers often aggregate the routing information of several adjacent ASs in order to contain the growth of their routing tables [Chen99]. We have therefore made use of the second method, that of router-level path traces.

A number of research and commercial groups regularly perform measurements on the Internet. Many of these measurements focus on just one application of the Internet, the world wide web (WWW), but there are also a number of projects targeting the Internet infrastructure, such as the Pinger project [PingER], CAIDA's Skitter tool [CAIDA], the Surveyor project [ANS] and the Test-Traffic measurements (RIPE-TT) of RIPE (Réseaux IP Européens), the collaborative community of organisations operating wide area IP networks in Europe and beyond [RIPE]. For our analysis we have used the data from the latter project. RIPE has placed a number of test boxes (around 70 in May 2002) near border routers of several ISPs, mostly in Europe, but also in the US and in the Asia-Pacific region. These test boxes send each other test messages, of which the test boxes record the route the message has taken through the Internet and the time needed for this trip. Because each packet travels through many intermediary networks on its way from source to destination, a fairly good view of the core of the Internet can be obtained.

3.2. Data Analysis

The raw data that the RIPE-TT project provides consists of two parts: the first part is *topology* data, consisting of the routes that packets took between test boxes at a certain time (called a *trace*). Each trace is consists of a list of typically 15 IP addresses, representing the routers that the test packet passed on its way. The second part of the data, the one way delay data, consists of the time (in milliseconds) that a test packet needed for its trip from one test box to another at a certain moment in time. By combining all trace data in a single graph, and by mapping IP addresses to ASs, a view of the topology of the Internet at the AS level can be obtained. Used the obtained topology, the end-to-end delay data can be attributed to inter-AS links, giving an indication of the performance of a certain inter-AS connection at a given time. This analysis is performed by a computer program developed for this purpose [Traces].

4. Results

By performing the data analysis described in the previous section on the RIPE-TT data for June 25, 2001 at 20:00 PM, Figure 3 is obtained. In this graph, the calculated delay for every inter-AS link is coloured from green (meaning no measurable delay) to red (up to 7500 ms). The nodes in the figure represent ASs. Due to the rules of use of the RIPE-TT data, the ASs are presented in an anonymous way, but all major ISPs (like UUNet, Cable&Wireless, Qwest, etc.) are present in the graph. The nodes in the graph with a large number of connections all represent one of the major ISPs. The same holds for the ASs that consist of the largest number of routers.

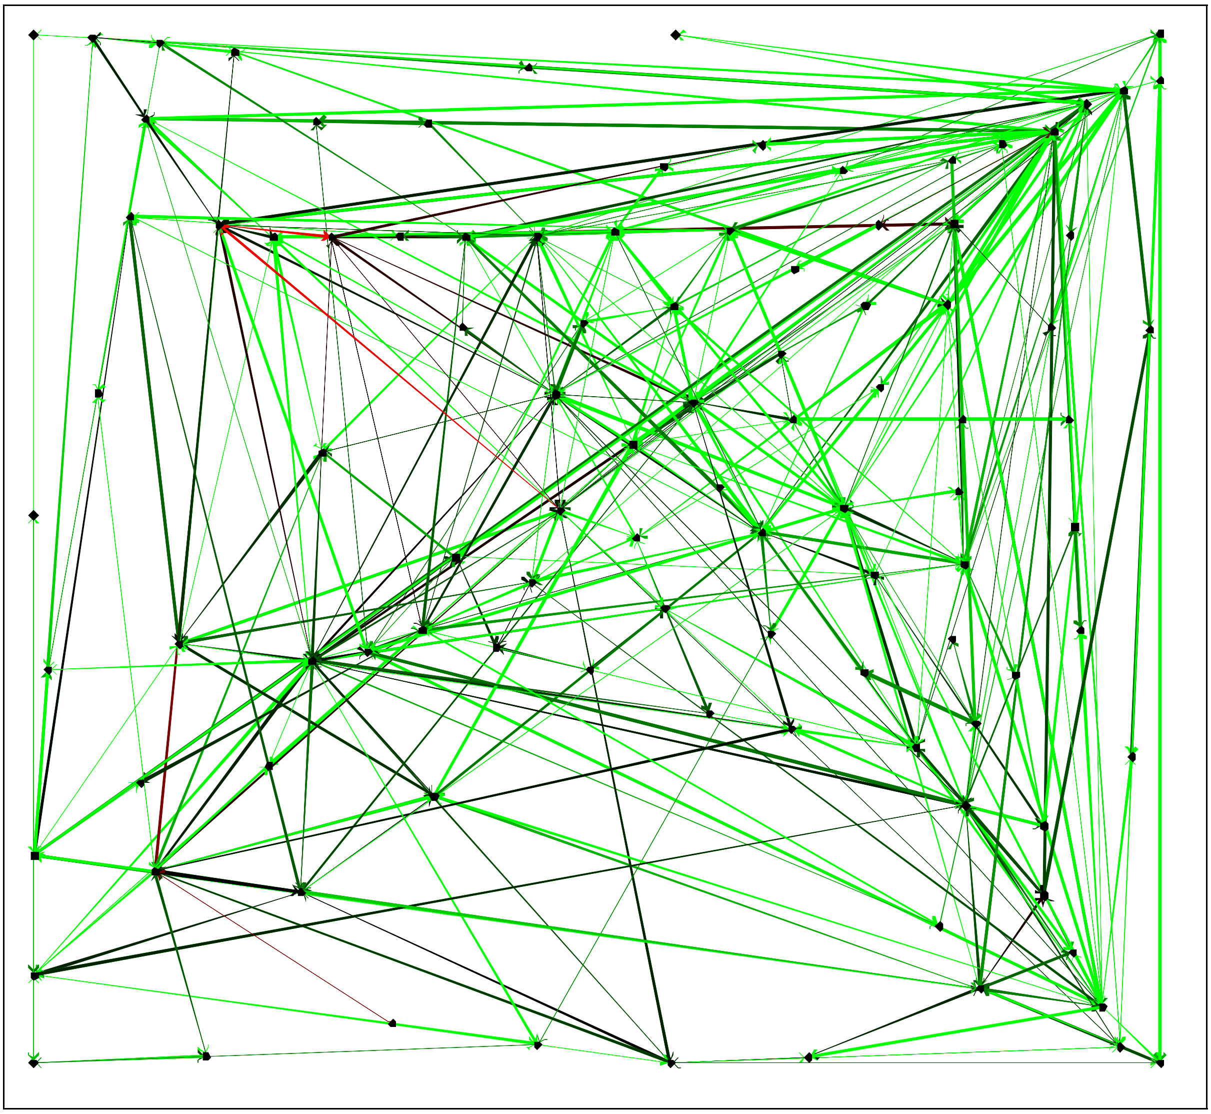

Figure 3: AS-level graph as measured by RIPE-TT at 25th June 2001, 20:00

5. Discussion

We have developed a method to create a map of the Internet on the AS level. This map can also be viewed as a map of the interconnection agreements between ISPs, measured from the outside, without asking each and every ISP what other ISP they interconnect with. When viewed over time, the maps can be used to assess current trends in the Internet infrastructure regarding issues like market power, concentration, dependence upon one or few parties for Internet service, and so on.

Future research will address the policy relevance of the AS-level Internet map, especially regarding its dynamics. The first results [Best2002] showed a remarkable geographical cohesion in the core of the European part of the Internet, indicating that interconnection relations between ISPs often coincide with geographical and cultural similarity.

References

[ANS] Advanced Network & Services, *Surveyor Project*, http://www.advanced.org/surveyor/

[Best2002] Jan-Pascal van Best and Willem G. Vree, “Locality of Internet Connections”, INET 2002, June 2002.

[CAIDA] Cooperative Association for Internet Data Analysis (CAIDA), *skitter*, http://www.caida.org/tools/measurement/skitter/

[Chen99] E. Chen and J. Stewart, *RFC 2519: A framework for inter-domain route aggregation*, February 1999

[Kende00] Michael Kende, *The Digital Handshake: Connecting Internet Backbones*, Federal Communications Commission, Office of Plans and Policy, Working Paper #32, September 2000, http://www.fcc.gov/opp/workingp.html

[Magoni01] Damien Magoni and Jean Jacques Pansiot, “Analysis of the Autonomous System Network Topology”, in *Computer Communication Review 31*, Number 3, ACM SIGCOMM, July 2001.

[Norton01] William B. Norton, *Internet Service Providers and Peering*, Equinix, 2001, http://www.equinix.com/press/whtppr.htm

[PingER] Internet End-to-end Performance Monitoring group, *The PingER project*, http://www-iepm.slac.stanford.edu/pinger/

[RIPE] RIPE NCC, *Test Traffic Measurements*, http://www.ripe.net/test-traffic/

[Tanenbaum97] Andrew S. Tanenbaum, *Computer Networks*, Academic Service, 1997

[Traces] Jan-Pascal van Best, *Traces software package description*, http://www.tbm.tudelft.nl/webstaf/janb/software