

Soundness of Chilean Internet Routes

Pablo Sepúlveda, Victor Ramiro, Tomás Barros
NIC Chile Research Labs, Universidad de Chile
Miraflones 222, Piso 14, 832-0198, Santiago, Chile.
{psepulv,vramiro,tbarros}@niclabs.cl

José Miguel Piquer
DCC, Universidad de Chile
Av. Blanco Encalada 2120, Santiago, Chile.
jpiquer@dcc.uchile.cl

Abstract—Internet is known as the net of networks. It is a huge decentralized network composed by smaller interconnected networks named autonomous systems. The interconnection of these autonomous systems is done by the Border Gateway Protocol (BGP). Several tools have been proposed to monitor BGP problems. However there is no simple indicator to rank Internet routes health which could be aware of BGP problems on long term periods. We want to answer: Is it possible to quantify the reachability from several places in the world to my network? In this sense, the major contributions of this work are: We present a new model to measure network soundness, specifically quantifying the routing reachability to a given network. The output is the percentage of the network that is present in the routing tables of selected routers around the world. We develop a systematic monitoring system based on the analytic model presented. Specifically we monitor all Chilean networks. Finally, We analyze the impact of the Chilean 2010 earthquake on the national network infrastructure, based on our model.

Keywords-Internet, Networks, Measurement, Routing, BGP

I. INTRODUCTION

Internet is a decentralized network composed by smaller interconnected networks named autonomous systems. The interconnection of these autonomous systems is done by the Border Gateway Protocol (BGP) [1]. BGP is known by its simplicity: it is based on sharing one's routing table among one's neighbours. For each new route received, the router will add the route to the routing table. If the new route is better than its previous best route, the router will share it with its neighbours. This sharing scheme is known as BGP convergence. Nevertheless BGP simplicity is also a threat for network security [2]: the sharing criteria is based on trusting neighbours. It is easy to break a router's chain of trust and spy data, hijack a network or provoke a denial of service attack.

To connect to a network through BGP, a router maintains in its routing table the sequence of autonomous systems the packet needs to traverse to reach a network, this is known as AS Path. A router maintains a set of networks and their corresponding AS Paths. The communication among neighbours is made through updates, there are two significant updates a neighbour can send: (i) an announce of a new path or (ii) a withdrawal of an existing path. When a neighbour sends an update for a new path, the router examines the

path and if it is better than its currently used path, the router replaces it on its routing table, announces the new path to its neighbours (appending its own AS number) and withdraws the older path. When the router receives a withdraw update and he was using that path, the router removes it from his routing table, and announce the withdraw to his neighbours. Then he chooses the next best path; if there is one, he promotes it to its routing table and announces it; if there is not, that network is now unreachable from this router.

In a normal day, every network used must be visible from all over the Internet. When a problem occurs in a point among the path between two networks, BGP convergence should look for an alternative route. For instance, on [3] there is an analysis of the effects of a fibre cable cut, measuring the changes in connectivity on the Internet when prefixes and routes change.

It can be critical for a organization to have a measure that rates how different clients around the world reach their services. Let us consider a organization that manages an IPv4 network. It publishes two subsets of the network by means of two different Internet service providers and keeps some IPs for future use. If it is critical for this organization to deliver services worldwide, one needs to know if its network is accessible anywhere even though it is published by different Internet service providers. This may sound obvious, but imagine that the web server is in one subnetwork and the database server is in the other subnetwork. If there is a problem in the BGP routes, it is possible that some clients are able to reach the web server but not the database server, whereas some other clients are able to reach the entire network (or more precisely the published network).

Several tools have been proposed to monitor BGP problems [4], [5], [6]. However there is no simple indicator to rank Internet routes health that is aware of BGP problems on long term periods. We want to answer: Is it possible to quantify the reachability from several places in the world to my network? In this sense, the major contributions of this paper are:

- 1) We present a new model to measure network soundness.
- 2) We develop a systematic monitoring system based on the analytic model presented. Specifically we monitor

all Chilean networks.

- 3) We analyze the impact of the Chilean 2010 earthquake on the national network infrastructure, based on our model.

II. A NEW METRIC FOR NETWORK SOUNDNESS

We propose a model for quantifying the routing reachability to a given network. The output is the percentage of the network that is present in the routing tables of selected routers around the world. Before going into the details of the model we present a basic notation and terms definitions to set a common language, then we explain the main issues related to measuring the degree of reachability, and finally we present the analytical model supported by a simple case of analysis.

A. Basic model conventions

In this paper, we consider a network as a finite set of IP addresses. We work with IP version 4, which defines addresses as 32 bits length binary number, written, as usual, into four octets separated by dots (A.B.C.D). The set of IP addresses is defined by a prefix p (an IP address) and a network mask m , the mask is the number of shared initial bits of the set, counting from the most significant bit of the address. Given this, we write a network $n(p, m)$ as follows:

$$n(p, m) = \{ip \in \mathbb{N} : ip \in [p, p + 2^{32-m}]\}$$

Networks fulfill several properties. One in our interest is the size of a network, which depends only on the mask m :

$$|n(p, m)| = 2^{32-m}$$

We define a subset N of networks that are reachable by Internet routers from the complete routable Internet (we denote this set as \mathcal{I}). This set is composed by networks that fulfill a given criteria. In our case, we define this criteria as: “the network is a Chilean network”, being defined taking as input the LACNIC network assignation table [7]. This is:

$$N = \{n(p, m) \in \mathcal{I} : n(p, m) \text{ is a Chilean network}\}$$

It is easy to see that every two networks in the set of networks N are one of two things: (i) one is a subnetwork of the other, or (ii) they are disjoint.

$$\forall n_i, n_j \in N \implies n_i \subseteq n_j \vee n_j \subseteq n_i \vee (n_i \cap n_j = \emptyset)$$

Finally, a *Collector* gathers routing tables from default-less routers. A default-less router has the routing table of all Internet. A collector table r maintains the set of networks this router can reach, this is:

$$r_j = \{n_i \in N : \text{collector } j \text{ can reach } n_i\}$$

To quantify the routing reachability to a network, we collect the BGP routing tables from a set of *Collectors* distributed around the world. These tables are filled with entries that relate a network $n(p, m)$ with the autonomous systems path to reach from the collector to the network.

$$n(p, m) : AS_1 AS_2 AS_3 \dots AS_n$$

In this example, the autonomous system from the collector router is AS_1 and the destination autonomous system is AS_n which represent the network $n(p, m)$.

B. Problem statement

Consider the Figure 1, which represents the network: $n(p, m) = 10.0.0.0/24$ ¹. Following the Classless Internet Domain Routing (CIDR) scheme, this network can be divided in halves, represented by two sub-networks with a bigger mask: $n_1(p_1, m + 1) = 10.0.0.0/25$ and $n_2(p_2, m + 1) = 10.0.0.128/25$.




Figure 1. Network tree for $n_1 : 10.0.0.0/24$ network. Showing two subnetworks A: $10.0.0.0/25$ and B: $10.0.0.128/25$

For instance, the owner of organization n wants to publish the entire network n to the Internet. We would like to know if it is always possible to find a path from every Internet router in the world to the destination router representing the network n . Looking at the entries for the network n at the routers, we can define a notion of reachability. Given a sample set of routers around the world, we can define the percentage of the network n which is actually present in the routing tables.

Even though, some problems are present in this approach that lead to different scenarios at a collector router. We could see the following entries:

- n : The whole network is seen at the routing table.
- n_1 **and** n_2 : The whole network is seen at the routing table, though this is inferred by seeing both sub-networks.
- n_1 **or** n_2 : Only a part of the network is seen at the routing table, since one subnetwork entry is missing
- no entry: The whole network cannot be seen at the routing table.

Which entries are seen can be useful to determine the healthiness of the network. For instance, seeing only n_1 or n_2 , i.e. a part of the network, can be due to a misconfiguration problem. However, it is a common practice that a

¹We use private Internet address just to illustrate the example.

organization reserves some network space for future use, which boils down into publishing a smaller subnetwork from a network, i.e. only n_1 or n_2 .

We can also use these entries to detect connectivity problems. Let a company have two links with two different Internet service providers, with each part of the network (n_1 and n_2) being published by a different provider. If one of the links fails, one subnetwork will be deleted from Internet for a period of time. Therefore, in order to detect such problems, one must analyze the routing tables of a collector over a long(er) period of time, moreover, one must analyze the coherence of the routing tables of multiple collectors. Note that this coherence must be some sort of equivalence in the sense that if one collector sees n and another one sees n_1 and n_2 , then both collectors see the same published network.

We introduce two basic set notions to rank the healthiness of the networks: (i) *announced* and (ii) *visibility*. As its name suggests, the *announced* indicator measures the percentage of the network that has been published to the Internet by the owner of the network. As we presented before, problems w.r.t. published networks can occur. The *visibility* indicator measures the percentage of the network that is actually present in a group of routing tables. This indicator establishes the healthiness of the network contrasting the values of the *announced* percentage of the network over the percentage that is actually present in the routing tables.

In order to define a metric of healthiness of the network, these indicators should fulfill two basic criteria:

- **(R1)** We need to collect data spread around the world to provide unbiased indicators, and
- **(R2)** We need temporal indicators that allow us to compare and rank over time the healthiness of the networks.

C. Abstract model framework

We present a model based on simple set theory and operations. Our model works with a set of disjoint networks $N = \{n_1, n_2, \dots, n_i\}$ and a set of collector router tables $R = \{r_1, r_2, \dots, r_k\}$. As usual, $|R|$ denotes the number of collectors in the R set. For the sake of notation simplicity we say that $n \in R \iff n \in \cup_{r_j \in R} r_j$. Given this basic basis, we define:

1) *Individual indicators*: We present a set of individual indicators defined over a network $n \in N$. However, they only give information about the single network n and its subnetworks if they are present. All the indicators are defined as a function ranged over $[0, 1]$. We will use indistinctly these values as numbers in $[0, 1]$ or as percentages in $[0\%, 100\%]$.

Presence: We define the presence $P(n, r)$ as the percentage of the network n found in the routing table of collector r . Values of $P(n, r) \in [0, 1]$, where 0 means that n nor any subnetwork is present, 1 means the whole network is present, and a number in between means that some of its subnetworks are present.

$$P(n, r) = \frac{|n \cap \cup_{n_i \in r} n_i|}{|n|}$$

The normal expected value for the presence $P(n, r)$ is the same as the published network, so if you publish half of your network, the presence should be 0.5. This indicator is used to construct the *Announced* and *Visibility* indicators.

Announced: We say a network is *announced* if it is present in any routing table r . We define $A(n)$ as follows:

$$A(n) = \frac{|n \cap \cup_{n_i \in R} n_i|}{|n|}$$

This definition shows intuitively that values of $A(n)$ indicate the percentage of the network that should be present in all routing tables. If the complete network n has been published, then the expected value for $A(n)$ is 1. Even though this number could be misleading since it suffices that $P(n, r^*) = 1$, for any r^* , to have $A(n)$ to be 1. In other words, if you publish a network n then an entry with this network should exist in at least one collector r^* .

If we do not often change the information of published networks, then this indicator should be stable over time. This property can be used to detect problems with BGP routing.

Visibility: We define the *visibility* $V(n)$ over the *announced* networks (if a network is not *announced*, we say that is not visible). Hence, the *visibility* of an *announced* network n is the average of the presence of the *announced* n in all collector routers r .

$$V(n) = \frac{\sum_{r_j \in R} P(n, r_j)}{|R|A(n)}$$

Notice that $V(n) \in [0, 1]$ should always be 1 when routing tables are coherent, meaning the whole amount of *announced* network is reachable from all collectors $r_j \in R$. Any number lower than 1 must be a warning of routing misconfiguration, because there are *announced* networks not visible all over the Internet.

2) *Group indicators*: Since several institutions manage more than one network, we define a group view of the indicators defined before. Here, *announced* and *visibility* indicators show the global healthiness of a set of networks.

Group Announced: Given a group of networks $G \subseteq N$, we define the *announced* of the group as the weighted ratio between *announced* networks and the size of all networks of the group. The weight function is the size of each network: $|n_i|$.

$$A_G = \frac{\sum_{n_i \in G} |n_i| A(n_i)}{\sum_{n_i \in G} |n_i|}$$

Group Visibility: Given a group of networks $G \subseteq N$, we define the *visibility* of the group as the weighted ratio between visible networks and *announced* networks for all networks of the group. The weight function is the size of the *announced* network: $|n_i|A(n_i)$.

$$V_G = \frac{\sum_{n_i \in G} V(n_i) |n_i| A(n_i)}{\sum_{n_i \in G} |n_i| A(n_i)}$$

D. Analysis of the model

The set of indicators defined by the model should be used in conjunction to be correctly interpreted. For the sake of comprehension, the following of this section will explain this through examples.

Let us take for example the Figure 2 to show different cases of analysis. In this example we shall just consider the individual indicators defined in Section II-C. This Figure shows a network $n(p, m)$ and its two subnetworks $n_1(p_1, m + 1)$ and $n_2(p_2, m + 1)$. We have information of three collector routers: $R = \{r_1, r_2, r_3\}$ and $|R| = 3$. Each box shows in green the portion of the network present in its routing table.

Scenario a: In Subfigure 2(a) we analyze the case when the complete network n is published, hence we expect to have $A_a(n) = 1$. We see that the collector r_1 contains an entry for the subnetworks n_1 and n_2 . The collector r_2 contains an entry for the complete network n and the collector r_3 contains an entry for the subnetwork n_1 . The presence of the network n in each collector is:

$$P_a(n, r_1) = 1, P_a(n, r_2) = 1, P_a(n, r_3) = 1/2 = 0.5$$

Since $P(n, r_1) = 1$ we say that $A_a(n) = 1$, just as we were expecting. However, it is easy to see that something is wrong about the entries in each collector: they are not the same. This means that each collector does not know the same amount of the network, hence it is not possible to reach the same contents from each collector in the network n . This is reflected by the *visibility* indicator since it is lower than 100%.

$$V_a(n) = \frac{1 + 1 + 0.5}{1 \cdot 3} = 83,3\%$$

Scenario b: Another scenario is depicted in Subfigure 2(b), in which the complete network n is published. In this case we see that the collector r_1 only contains an entry for the subnetwork n_1 . The collector r_2 only contains an entry for the subnetwork n_2 and the collector r_3 does not contain any entry. We compute the presence of the network n on each collector:

$$P_b(n, r_1) = 0.5, P_b(n, r_2) = 0.5, P_b(n, r_3) = 0$$

Notice that collectors r_1 and r_2 contain two different subnetworks, but they form in fact the n network. Hence,

we say that the $A_b(n) = 1$ (equivalent to say that the 100% of the network n is *announced*). In this case it is also clear that there are problems with the information spread among the different collectors, which is reflected by the *visibility* indicator being less than 100%.

$$V_b(n) = \frac{0.5 + 0.5 + 0}{1 \cdot 3} = 33,3\%$$

As we can see $V_a(n)$ is greater than $V_b(n)$. This is obvious since in the scenario (a) just the collector r_3 has missing information in its routing table. In case (b) the three collectors have incomplete data.

Scenario c: Finally, we analyze the case when just a subnetwork of n is published on Internet, depicted in Subfigure 2(c). We see that the collectors r_1 and r_2 contain an entry for n_1 and collector r_3 does not contain any entry. Notice that collector r_1 and r_2 contain now the same network, which is in fact a subnetwork of the n network. The presence of the network n in each collector is:

$$P_c(n, r_1) = 0.5, P_c(n, r_2) = 0.5, P_c(n, r_3) = 0$$

Hence, we have that $A_c(n) = 0.5$ (equivalent to say that the half of the network n is *announced*). Now the *visibility* of the network n is different:

$$V_c(n) = \frac{0.5 + 0.5 + 0}{0.5 \cdot 3} = 66,6\%$$

In this case we see that $V_c(n)$ is greater than $V_b(n)$. This is because in case (c) we published a smaller amount of the network. Therefore, collectors r_1 and r_2 have correct data and just r_3 has wrong data.


As we see with these three scenarios, the analysis of the indicators must be done in pairs: each indicator gives us a partial information of the status of the health of the analyzed network.

In conclusion a quick view at different values for the indicators have different meanings: (i) If *announced* is constant on time and *visibility* = 1, the network is healthy, (ii) if *visibility* < 1, the network have some kind of problem of misconfiguration, (iii) if *announced* increases, it means more of the network is being published, but if it decreases for a period of time, it probably means there is some kind of problem, like an outage or a misconfiguration.


III. IMPLEMENTING THE MODEL

Given the model presented in Section II, we define the steps necessary to implement and calculate the model, and present the input and the output expected on each step. We also show how this implementation meets the requirements defined in II-B.


This model is implemented based on the RIPE-RIS collectors project [8]; RIS is a RIPE NCC project that collects and stores Internet routing data from several locations around the



(a) Case where the whole network is *announced*, but it's not seen in all collectors



(b) Case where the whole network is *announced*, but it's not seen in all collectors



(c) Case where a subnetwork is *announced*, but it's not seen in all collectors

Figure 2. Figure to show the difference between Announced and Visibility in the model

globe. RIS offers this data to the Internet community; this data is particularly useful to our investigation.

Initially, we started our database with the set of networks we called super networks (SN), which have two properties:

- 1) Each network is disjoint to the other networks

$$n_i, n_j \in SN, n_i \neq n_j \implies n_i \cap n_j = \emptyset$$

- 2) The union of the networks in the set covers all the interesting networks, the set N . This is, SN is a partition of the interesting network space

$$ip \in n \in N \implies ip \in \cup_{n_i \in SP} n_i$$

Then we define a four-stages process that is repeated automatically everyday:

- *Data Fetching*: We automatically download the data each day from the 12 RIPE collectors, obtaining the raw data; this data is in a format known as MRT [9] and we use libbgpdump [10], a library designed to read the data. The data we use is the entire BGP routing table for each collector, in the same form that the router received it, i.e. the announcing updates. The information contained in the updates are: the time it was sent, the prefix, the IP and the AS number of the sender,

and the BGP attributes. Among the attributes is the AS path, which is the list of autonomous systems the packet needs to traverse to reach the network; the next hop, which is the next router which will forward the packet; the local preference, which is used to choose between various paths; and the communities attribute which is used to transmit additional data for extensions.

- *Data Analysis*: On the second step we filter the data by the networks we are interested in, the N set; reducing all the information obtained to the minimum we require, a list of networks for each collector. In this list are all the interesting networks, which this collector can reach. To properly filter the data, we need to consider that subnetworks of networks in the N set are also relevant for our analysis.
- *Model Calculus*: We then read the list and apply our model, i.e. we calculate the *announced* and *visibility* values for each of the networks we are interested in, i.e. networks in the N set. To calculate these values, we developed a network-tree structure to represent the network. In this structure we mark each network we see, and recalculate recursively the value of *announced* or *visibility* as the case. We group each tree in a forest structure to search efficiently in which network-tree we

need to mark the networks we see.

- *Data publishing*: Finally, we publish the data, the *announced* and *visibility* values for the networks on the Internet so anyone can see the metrics for any network for a given day, and we support the data with graphics so its easy to visualize changes and spot the most significant events.

IV. MODEL APPLICATIONS

In this section we present two different implementations of the model following the guidelines defined in Section III.

A. Monitor of Chilean networks

We now develop a quality monitor of Internet routes for Chilean networks based on the LACNIC whois information [7]. We define the Chilean supernetworks as the set of networks that define the partition of the Chilean network space. We analyze the collector information from RIPE-RIS project [8]. This project provides information from 12 collectors distributed around the world. We study 416 companies with IPv4 networks, representing a total of 682 supernetworks.

Daily we calculate the *announced* and the *visibility* for all Chilean networks. We define two groups indexes based on the networks origin:

- 1) Company indicator: We applied the group indicator defined in Section II-C2 to calculate aggregated values for each company with Chilean networks. The information is based on the LACNIC whois table [7]. We obtained values for *announced* and *visibility* for each company, allowing us to compare them. All the companies but one had 100% visibility, the only one not having a complete *visibility* had around 99% visibility. This means that the chilean networks are very sound, and only one company have problems on its configuration.
- 2) Chilean networks indicator: we also applied the group indicators to all the Chilean networks as a whole. We obtained values for *announced* and *visibility* for the Chilean networks as a whole, which give us a good measure of the Chilean networks healthy.

In the case of Chilean networks, we see that normally the 80% of the assigned networks are *announced* and the *visibility* of those networks is near 99%.

In the monitor we display all the data for a given day, as well as the aggregated data. The monitor plots the aggregated data for each company and for the whole Chilean networks. The monitor also shows in a map chart the *visibility* for each collector of chilean networks. This monitor is available at [11].

B. 27-F Earthquake

The 2010 Chilean earthquake occurred off the coast of the Maule Region of Chile on February 27, 2010, at 03:34

local time (06:34 UTC), rating a magnitude of 8.8 on the moment magnitude scale. This earthquake affected the Internet connectivity in an unknown scale. As an interesting case of study we tested our Chilean network monitor to get the *announced* and *visibility* indicators the day before and after of the 27-February earthquake. In Table I we see more details of the normal values for the indicators; for this, we used the 25 of February as a normal day to calculate the variation the indicators. We see that usually 82% of the Chilean assigned networks are *announced* to Internet and those networks are highly visible, with a *visibility* value of 100%. However, on 27-February we see how the *announced* indicators drops down to 31,10%, which highlights that a significant number of networks stopped being *announced* to Internet. This shows us at a glance, based on the Chilean Internet monitor measures, that 64,45% of the Chilean networks were unreachable from outside of Chile². Obviously this was caused by the earthquake and not a misconfiguration issue. Even though, from the 31,10% *announced* that day we see that it was highly accessible with a 96,43% of *visibility* indicator.

In normal operation, the monitor of the Chilean Internet takes one snapshot from each collector per day, specifically at 8:00 UTC. This was not enough to measure the actual impact of the earthquake on the Internet infrastructure during the 27-February and over the following days since it is too coarse to detect fast changes in routing tables. Therefore, we have extended the monitor to process the collectors data and compute all the indicators to determine the minute by minute information of the networks.

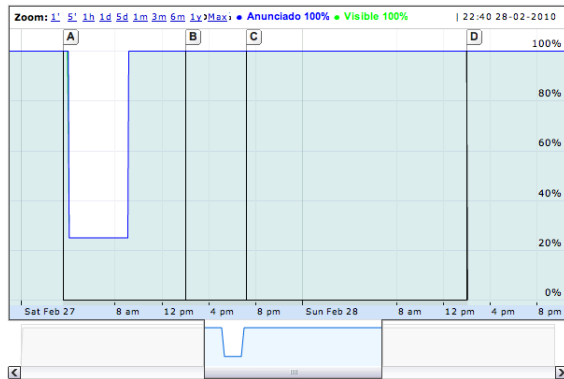
The second analysis showed that Chilean Internet had a massive failure after the quake between 4:00 and 12:00 local time. In Figure 3, we present three anonymous Chilean networks to contrast the damage suffered by them. The time window is set between 00:00 of 27-February and 22:40 of 28-February. The figure depicts both indicators: the *announced*, by a green area, and the *visibility*, by a blue area. In these figures we see that both areas are indistinguishable one of the other. This is explained by the high value of the *visibility*, which is always around 100%. There are also major milestones marked in the figures: (A) Earthquake at 3:34:17 2010-2-27 and (B), (C) and (D) show the recovery of the power supply w.r.t. a normal day in the central interconnected energy system of Chile. They show 24%, 30% and 43% respectively of the energy recovery.

It is clear that in the network depicted by Figure 3(a) there is a partial outage: a small amount of the network is still *announced* after the earthquake. This can be explained by a company having more than one Internet service provider, where one failed and the other survived to the earthquake. In Figures 3(b) and 3(c) we see a complete outage of the


²Since we do not count with a collector in Chile, the behaviour of the Chilean networks within Chile cannot be measured during the earthquake.

Dates	Announced	Visibility	Visible w/r a normal day	Not Visible w/r a normal day
25-Feb 2010	82,05%	100,00%	100,00%	0,00%
26-Feb 2010	81,88%	100,00%	99,79,%	0,21%
27-Feb 2010	31,10%	96,43%	36,55%	64,45%
28-Feb 2010	78,40%	99,96%	95,51%	4,49%
01-Mar 2010	79,53%	100,00%	96,90%	3,10%


Table I
CHILEAN EARTHQUAKE INTERNET SHORTAGE



(a) Partial short outage



(b) Complete short outage



(c) Complete long outage

Figure 3. Indicators over time in the earthquake analysis

network. In the first case we see a recovery around 12:00 local time, which is correlated with the recovery of the power supply in Santiago. Probably they did not count with enough power supply to continue on service for a long shortage. In the later we see a bigger impact, which lasts throughout the day.

We summarized the data of all Chilean supernetworks outage to quantify the complete impact. For each supernetwork we calculated the outage time. Then we counted the number of networks that had an impact of over 30% of their *announced* networks. We see that 29,63% of the supernetworks, this is 96 Chilean supernetworks, were affected by less than 30% of their *announced* networks, hence we say they were unaffected. The remaining 70,37% supernetworks, representing 228 Chilean supernetworks, were unavailable signaling a major flaw in the Chilean networks. In Figure 4 we present a histogram of the outage time versus the total number of networks that were unavailable. The histogram shows two different groups of supernetworks with respect to the outage time distribution. The first one representing the 54,32%, this is 176 supernetworks, follows a normal distribution with mean 9 hours and variance around 5 hours. These are the networks that presented problems but quickly recovered normal operation. The second group are supernetworks that were out of service for more than 24 hours; this group corresponds to 16,05% of the supernetworks, representing a total 52 Chilean supernetworks.

This outage has been confirmed by the Google traffic transparency report [12]. This report provides information about traffic to Google services around the world. Each report shows historic traffic patterns for a given country and service. Graphs are normalized and scaled in units of 0 to 100. Figure 5 shows the YouTube traffic from Chile. The average traffic is 60,00 points in a normalized scale, though on 27-February this index fall 19,69 points.

More information of the earthquake impact on the Chilean networks can be found at [13]. A special video to visualize the impact of the earthquake during the 27-February is available ³.

V. RELATED WORK

The first study about Internet health was done in 1997 [14]. The authors observed for nine months the traffic

³<http://www.youtube.com/watch?v=ZOTjnHXRai4>




Figure 4. Histogram showing the number of networks with an outage of X hours



Figure 5. Google YouTube Traffic report November 2009 - May 2010, in a normalized scale

between five BGP messages interchange points. The main results were the large amount of updates done, most of them bogus or redundant. Almost ten years later, a similar study [15] obtained data from August 2005 to January 2006, showing a substantial improvement on BGP health; the bogus messages were significantly less, and only redundant messages remained similar.

A specific study of a cut in the Mediterranean fibre cable cut was done by RIPE [3], measuring the impact of the cut on the healthy of the network.

On those studies, the “metric” was the amount of bogus or redundant updates. This metric is significant when measuring the impact of the problem in the whole Internet, but is not that useful if one wants to compare the healthiness between two networks, two autonomous systems or the same network on two consecutive days, or compare two consecutive months.

Renesys studied the Internet the day the World Trade Center collapsed, September 11, 2001 [16]. In their analysis they used as metric the number of reachable prefixes, and the number of announcements to measure the instability. They concluded the Internet survived thanks to sufficient redundancy. Our work creates a better metric, because it includes the size of the network for a more precise understanding. In our study we have concluded that we did not have enough redundancy so our Internet did not survive as in their study.

Besides those studies, there are initiatives to store the routing information and make that data publicly available. Among them there are the University of Oregon Route Views Project [17], the RIS project from RIPE [8] and the BGP monitor from the MIT [18].

Some initiatives use this data to build systems to monitor the BGP routes, among them there are BGPMon [5] and Cyclops [4], which allow users to query and send alerts when something changes in order to detect malicious redirection or errors caused by a misconfiguration. This monitor systems throw alerts when changes occur in each route, or abnormal behaviour of autonomous systems, such as bogus announcements or new peerings.

The alarms are significant against attacks, but do not measure the reachability of the network. Our system provides a measure to compare different networks, besides it can be used to group networks of the same company or country allowing further comparison. Our system can detect problems in reachability caused by power outages, fiber cuts or many other problems, though it does not detect attacks as BGPMon or Cyclops.

VI. CONCLUSIONS

In this paper we have presented a simple model to measure network reachability based on a vision of networks as sets of IPs. We have introduced two indicators, *announced* and *visibility*, which represent the amount of network that is spread and routable from different points in the Internet backbone. These indicators help us to measure network healthiness.

Based on the LACNIC list of network, and the RIPE-RIS project, we have implemented the model and have built a monitor for Chilean networks. This monitor run the model once per day to rank Chilean networks healthiness over time. Based on this measures we found than around 99,99% of the Chilean networks present the expected values for the indicators. This means that the networks that are *announced*

are accessible from all the collectors routers selected for the study around the world. In other words, we can say that Chilean networks, and hence their contents, are accessible over the world.

We have also presented an extended tool to perform a forensic analysis of the consequences of the 27-February Earthquake over the Chilean Internet Infrastructure. This study has showed that 64,45% of the Chilean networks were not reachable that day. Also we have showed that the 70,37% of networks were in average 9 hours out of service.

Future work includes an extension to track and measure IPv6 networks. This extension is interesting for two reasons: (i) IPv6 is the next generation of the Internet protocol; and more importantly (ii), we can track and save the new Internet topology from its early stages. Also we plan to change the sampling frequency in order to obtain better metrics for smaller outages or problems. Finally, we want to install a RIPE-RIS collector router in a Chilean Internet exchange point to have a local view of the measures.

ACKNOWLEDGMENTS

The authors would like to thanks to Antonio Cansado, for his ideas to improve this work; all the NIC Chile Research Labs, team for the support and comments; and to the NIC Chile team, for their help on the development of this work.

REFERENCES

- [1] Y. Rekhter and T. Li, "A Border Gateway Protocol 4 (BGP-4)," RFC 1771 (Draft Standard), Internet Engineering Task Force, Mar. 1995, obsoleted by RFC 4271. [Online]. Available: <http://www.ietf.org/rfc/rfc1771.txt>
- [2] S. Murphy, "BGP Security Vulnerabilities Analysis," RFC 4272 (Informational), Internet Engineering Task Force, Jan. 2006. [Online]. Available: <http://www.ietf.org/rfc/rfc4272.txt>
- [3] R. Wilhelm and C. Buck, "Mediterranean fibre cable cut - a ripe ncc analysis," <http://www.ripe.net/projects/reports/2008cable-cut/index.html>.
- [4] R. V. Oliveira, M. Lad, and L. Zhang, "Cyclops," <http://cyclops.cs.ucla.edu>.
- [5] A. Toonk, "Bgpmon," <http://bgpmon.net/>.
- [6] L. Colitti, G. Di Battista, F. Mariani, M. Patrignani, and M. Pizzonia, "Visualizing interdomain routing with bgplay," *Journal of Graph Algorithms and Applications, Special Issue on the 2003 Symposium on Graph Drawing, GD '03*, vol. 9, no. 1, pp. 117–148, 2005.
- [7] Latin American and Caribbean Internet Addresses Registry (LACNIC), "Lacnic delegated resources," <ftp://ftp.lacnic.net/pub/stats/lacnic/delegated-lacnic-latest>.
- [8] RIPE, "Ripe-ncc. routing information service project," <http://www.ripenc.org/projects/ris/tools/index.html>.
- [9] L. Blunk, M. Karir, and C. Labovitz, "Mrt routing information export format," <http://tools.ietf.org/id/draft-ietf-grow-mrt-04.txt>.
- [10] D. Ardelean, "libbgpdump," <http://www.ris.ripe.net/source/>.
- [11] NIC Chile Research Labs, "Visibility of the chilean routes on internet," <http://www.niclabs.cl/visibilidad>.
- [12] Google, "Google transparency," <http://www.google.com/transparencyreport/traffic>.
- [13] NIC Chile Research Labs, "Internet and the 27f earthquake," <http://www.niclabs.cl/terremoto>.
- [14] C. Labovitz, G. R. Malan, and F. Jahanian, "Internet routing instability," *SIGCOMM Comput. Commun. Rev.*, vol. 27, no. 4, pp. 115–126, 1997.
- [15] J. Li, M. Guidero, Z. Wu, E. Purpus, and T. Ehrenkranz, "Bgp routing dynamics revisited," *SIGCOMM Comput. Commun. Rev.*, vol. 37, no. 2, pp. 5–16, 2007.
- [16] A. Ogielski and J. Cowie, "Internet routing behavior on 9/11," <http://www.renesis.com/tech/presentations/pdf/renesis-030502-NRC-911.pdf>.
- [17] U. of Oregon, "University of oregon route views archive project," <http://www.routeviews.org>.
- [18] N. Feamster, D. Andersen, H. Balakrishnan, and F. Kaashoek, "Bgp monitor," <http://bgp.lcs.mit.edu/>.