

# AS Path-Prepending in the Internet and Its Impact on Routing-Decisions

Software Development Project

Institute of Network Architecture  
Technical University of Munich  
85748 Garching bei München  
Prof. Anja Feldmann, PhD

Advisor: Wolfgang Mühlbauer

by  
Zhi Qi  
(ytqz@mytum.de)

Submit Date: October 20, 2006

# Abstract

The Internet is made up of more than 20,000 autonomous systems (ASes), each of which is a collection of IP networks and routers with a single and clearly defined routing policy. The Border Gateway Protocol (BGP) is the core routing protocol of the Internet. It works by maintaining a table of IP networks which designate network reachability between ASes.

Internet Service Providers (ISPs) use BGP to exchange information between their routers. They seek to build networks that will accommodate varying traffic loads, and be robust to link and router failures. For this aim, a better understanding of the BGP routing is very necessary.

AS-path prepending is a common method to influence BGP routing decision. By prepending the local ASN multiple times, ASes can make advertised AS paths appear artificially longer. AS path-prepend is widely applied in the Internet. It can be found in 15.4% AS paths we have observed. Analyzing data from BGP table dumps, we try to shed light on the importance of AS-path prepending in the Internet. Simple experiments are supposed to study how observed prependings can be integrated into interdomain-routing models of the Internet.

Abstract.....	2
1 Introduction.....	4
1.1 Outline.....	4
2 Background.....	6
2.1 BGP.....	6
2.1.1 BGP Basics.....	6
2.1.2 BGP Path Selection.....	7
2.2 Data Sets.....	8
2.3 C-BGP – A BGP Routing Solver.....	9
3 AS Path-Prepending in the Internet.....	10
3.1 Approach for Getting AS Path-Prepending.....	10
3.2 Basic Statistics on AS Path-Prepending.....	10
4 Evaluation of AS Path-Prepending Policy through Simulation.....	14
4.1 General Approach.....	14
4.2 Construction of AS-Level Connectivity Graph.....	15
4.3 Implementing Simulation.....	16
4.4 Obtaining Statistics of Simulated Routings.....	17
4.5 Analysis of Simulation Result.....	19
5 Conclusion and Future Work.....	23
6 List of Tables.....	24
7 List of Figures.....	25
8 Bibliography.....	26

# 1 Introduction

Today almost everybody uses the Internet on a daily basis. Email, searching eBay for a perfect gift, reading the news, or just “surfing”; it has become an integral part of our everyday lives.

The Internet is made up of more than 20,000 Autonomous Systems (AS). An AS is a collection of IP networks and routers which is administrated independently and has a single and clearly defined routing policy [WIKI01]. A unique AS Number (ASN) is allocated to each AS for use in Border Gateway Protocol (BGP) routing. With BGP, ASNs are important because they uniquely identify each network in the Internet [WIKI02]. BGP is the core routing protocol of the Internet. It works by maintaining a table of IP networks which designates network reachability between ASes [RFC4271].

Unfortunately the BGP routing decision is not completely based on system-wide criteria (e.g. shortest AS path). ISPs sometimes need to manually control the incoming and outgoing traffic in the interest of economic or performance. ISPs are also required to adjust their network to provide new services and enhance reliability and security. A better understanding of what is going on in the routing layer is therefore necessary.

C-BGP [QU01], a scalable BGP routing decision-process simulator developed at the Université Catholique de Louvain, is used to run large-scale simulation of BGP routing decision in our work. A model is designed and built, which is supposed to reproduce the inter-domain BGP routings (also referred to as AS paths). The model generates an AS connectivity graph relying on BGP table dumps derived from the Internet. By this model each AS is reproduced as a single router and direct connections between ASes (AS-level peerings) are obtained for further simulation.

On the basis of this model we want to reproduce BGP routing considering AS path-prepending policy. ISPs are sometimes using AS path-prepending to make certain routers less preferred by the BGP routing decision-process. Every time, a BGP message is forwarded to a router in a different AS, which has AS path-prepending policy for the specified BGP message, the ASN of the local AS is prepended several times according to the configuration of this policy before announcing to the next AS.

In our work we collect statistics of AS path-prepending from the Internet. On the basis of the statistics, we design a simulation model with consideration of path-prepending policy. The simulated result is later compared with the AS paths from the real Internet and the comparison is supposed to disclose the impact of AS path-prepending policy on the BGP routing decisions .

## 1.1 Outline

This work is organized as follows. We start with some background information on the BGP. After that we present data sets needed by the simulation and two different methods to get original BGP routings from the Internet. Next is an overview of BGP routing decision-process simulator C-BGP which is used in the later simulation.

Chapter 3 focuses on AS path-prepending policy used in the Internet. Here we firstly explain the approach to get statistics on AS path-prepending and analyze the result after that.

The aim of Chapter 4 is to present the exact approach of our simulation and the collected statistics. First of all is an overview of the simulation approach. Then it is described how the data sets derived from the real Internet are transformed into an AS-level connectivity graph. Next is the implementation of simulation. After that we see the comparison of the simulated routings and that of the real Internet. The result of this comparison is again compared with that of the simulation without considering prepending policy so that the impact of this policy on BGP routing decision can be revealed. At last we explain the reasons for the experiment result. Finally, we conclude our work and present future directions in Chapter 5.

## 2 Background

The first part of this chapter gives some fundamental information of BGP. Next is a brief introduction of the data sets used throughout the simulation. And then is an overview of the decision-process simulator C-BGP and its assistant application, which is designed to make large-scale simulation quick and convenient. The content of this chapter is the basis of the remaining parts of this work.

### 2.1 BGP

In the Internet, routers are used to direct the traffic from one point to another. Depending on how "distant" two points are, more than one router may handle the traffic, with each router or "hop" moving the traffic closer and closer to its target. In order to find the "next hop" router the router has a routing table which is usually maintained automatically using routing protocols. The routing protocols allow routers to exchange and share subnet location or topology automatically.

Inside a network, internal routing protocols such as RIP, EIGRP, OSPF, and IS-IS are used to dynamically advertise routes to route local traffic efficiently. Which internal protocol a site chooses depends on hardware support, staff experience, and internal policy. However, across the Internet between ASes, only one external routing protocol is used to exchange route information: BGP. BGP advertises routes of one subnet to other networks. [PA01] (Figure 2.1)

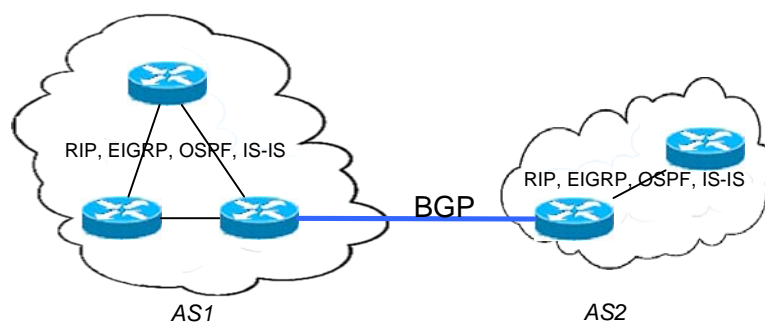


Figure 2.1 Internet routing protocols

#### 2.1.1 BGP Basics

As we've seen, BGP is designed to exchange routing information between ASes. Here's how BGP works:

In BGP, each AS has its own ASN, which is used to identify this AS to the world. ASNs are fundamental to BGP: Every time a route is advertised within BGP, the route is "stamped" with the ASN of the router doing the advertising (propagating). A sequence of ASNs forms an AS path, a kind of trail that literally reflects how the route became known to any router.

Next, just as the router stamps its ASN in the AS path when it advertises a route, each of the BGP peer routers prepends its ASN to the AS path, and in turn, advertises the route to each of its

external BGP peers. This propagating/stamping step continues as each AS prepends its ASN to the AS path and passes the route along. Table 2.1.1 shows an example of AS path.

1668 7018 2386
----------------

Table 2.1.1 An example of AS path

Obviously, not every BGP router is directly connected to every other BGP router. All routes in a BGP routing table are summaries of network connections somewhere out on the Internet. And AS paths reflect how each route is advertised, providing an insight into the routing in the Internet and could be used to determine the "shortest" AS path to a given target. [PA01]

In effect, BGP is like a road map for the Internet: BGP lists all of the ASes, records the major routes between those ASes, and select which paths provide the best "route" to the destinations. Once inside the network, internal routing protocols act as street maps to deliver traffic "the last mile" to a specific address. [WIKI03]

Most Internet users do not use BGP directly. However, since most ISPs must use BGP to establish routing between one another, it is one of the most important protocols of the Internet.

## 2.1.2 BGP Path Selection

It is possible to define a set of ingress filter rules for each BGP session. If a route is received over a BGP session, the router checks if it is in accordance with the ingress filter rules. If yes, the router may manipulate some attributes of the route before it is forwarded to the BGP decision-process.

The BGP decision-process is used to select a preferred route. In BGP "shortest" doesn't always mean best. Indeed, routes learned via BGP have also many associated properties that are used to determine the best route to a destination when multiple paths exist to this destination. These properties are referred to as BGP attributes. The AS path is just one of many BGP attributes that could decide the best routes. Policy routing (routing with BGP attributes) lets the network or system administrator override BGP's default behavior to choose "longer," but better performing paths. In fact, policy routing is a selection of criteria called the best path selection algorithm. These criteria are applied to multiple paths to the same destination in the specified order until only one path remains. (See [CR01] for more information)

If the new best route is different from the old one in the BGP routing table (referred to as Routing Information Base (RIB)), the RIB is updated.

Before being propagated to the neighbors, the new route may pass several egress filters which may also manipulate the BGP attributes of the route like ingress filters.

Among the BGP attributes there is an attribute called AS path. Shorter AS paths are preferred over longer AS paths. AS path-prepend is a common method to influence path selection. ASes not wanting a particular path used may prepend local ASNs in the AS paths advertised by them multiple times in order to make the AS path appear artificially long. [PA01] Take the AS path "1 5 2" as an example, if AS5 wants it less preferred, it may prepend its ASN multiple times (here 3 times) and the path is converted into "1 5 5 5 2". If there is another route "1 3 4 2" which was not selected as the best path because "1 3 4 2" had shorter AS, this route is now the better by comparing. (Figure 2.1.2)

According to the statistics on our data set (see next section), there exist path-prepend in 15.4% AS paths.

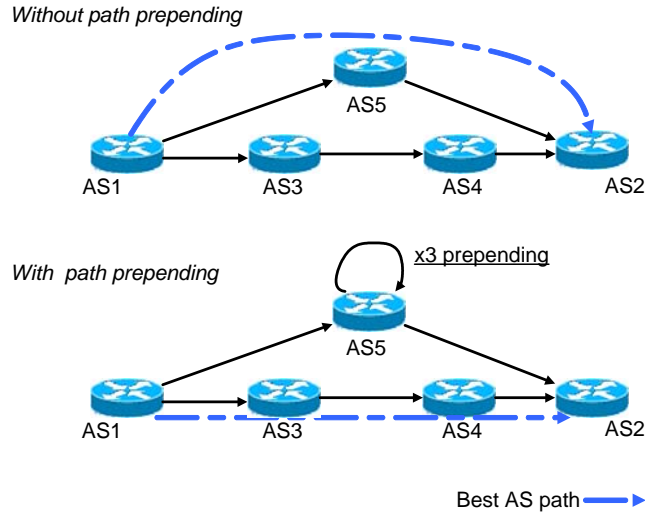


Figure 2.1.2 Effect of AS path-prepending on BGP routing selection

## 2.2 Data Sets

The recent progress in simulator makes it possible to simulate BGP routing decision-process on a large-scale. However, a set of basic data is a prerequisite for our work, upon which the simulation model can be designed. This section presents the data sets that are used throughout our simulation. There are two sources of BGP data sets: BGP table dumps and BGP updates. BGP table dumps provide a snapshot of the RIB whereas updates contain a time series of route changes. Concerning table dumps, it is sufficient to log on the router and query the RIB. BGP updates can be got in two ways: One can dump the IP packet trace on TCP port 179 which is used by the BGP sessions or one can use “route collector” software to collect BGP data on the real routers.

Our work relies on BGP routing table dumps from the Internet. All the data was recorded on 07/18/05 by using the “route collector” technique [MU01].

BGP table dumps can be inferred to an AS-level connectivity graph. The connectivity information fulfills two requests of our simulation: Firstly, it serves as the starting point of the simulator. The simulation is designed on the basis of the information. Secondly, it is used for the validation of the simulation by checking whether the simulator selects the same paths as those observed in the reality.

The used raw data provides an insight into the basic connectivity graph of the Internet. We extract unique AS paths from the data sets and discard private prefix announcements. From AS paths we can get AS-level peerings, which imply direct connections between two ASes. Some basic statistics are as follows:

distinct AS paths (ignoring AS prepending)	4,724,722
Number of ASes	21,173
AS-level connections (peerings)	62,595

Table 2.2 Basic statistics on used raw data (Excluding private ASes)

Table 2.2 shows that 21,173 ASes and 62,595 AS-level connections are identified from 4,724,722 unique AS paths obtained from our BGP collectors. Note that restricted by the number of

observation point, our data does not reflect the complete status quo of Internet at a fixed point in time. It could be possible to see two different best paths to the same destination in the case that there are several different BGP feeds from the same router.

## 2.3 C-BGP – A BGP Routing Solver

One goal of our work is to reproduce the AS paths. For this goal there are two prerequisites: The used data set which serves as feeding of the simulation and a simulator tool used to reproduce the AS paths. We have got the raw data set in the section 2.2. Concerning the simulator we select C-BGP which relies on an AS-level connectivity graph as an initial model and can reproduce BGP routing on a large scale.

There are surely also other simulators. The most closely related works are [FE01, FWR01]. The aim of [FE01] was to provide the networking industry with a software system to support traffic measurement and network modeling. This tool is able to model intra-domain routing and study the implications of local traffic changes, configuration and routing. However, [FE01] does not model the inter-domain routing protocol. [FWR01] proposes a BGP emulator that computes the outcome of the BGP route selection process for each router in a single AS. This tool does not model the flow of the BGP routes inside the AS, so it does not reproduce the route filtering process occurring within an AS. Neither of these tools is publicly available. [QU02]

C-BGP is an efficient BGP decision-process simulator developed at Université Catholique de Louvain (UCL), Belgium. C-BGP can be used to experiment with modified decision-processes and additional BGP attributes. It can also be used to evaluate the impact of input/output policies on the routing tables of other ASes. In our work, it takes the AS-level connectivity graph as input and outputs the routes selected to all possible destinations for each AS. C-BGP is able to consider all Internet domains (more than 20,000), provided that each is modeled with a single router.

The easy configuration of policies in C-BGP is a major advantage too. A configuration file with ingress and egress filter rules can be set up with Cisco-like commands. The simulator runs on the basis of this configuration file that comprises the used routers as well as ASes and reproduces the selected best routes after everything has converged. (See [QT01] for more information)

The drawback, however, is that dynamical aspects like BGP convergence times cannot be investigated with this simulator. But we are only interested in the state of the routing (routing tables) after everything has converged in our work. In this regard, the decision-process simulated by C-BGP is sufficient.

Because of the large amount of ASes and AS-level peerings, it is impossible to set up the configuration file for C-BGP by means of manual input. A tool that can automatically set up large-scale simulation on the basis of AS-level connectivity graph is needed. Fortunately an application is already developed to ease this work. The application takes AS-level connectivity graph as input and guides C-BGP to do the simulation so that output information is gathered and a summary is given.

The model can reproduce not only inter-domain but also intra-domain connectivity. Our work simplifies the BGP routing decision and only inter-domain is concerned. We use this model in a very simplified way by only considering inter-domain connectivity. (See [MU01] for more information)

## 3 AS Path-Prepending in the Internet

Among the policies that BGP uses to make routing decision, AS path-prepending policy is an important one. It is widely applied in the Internet, 15.4% AS paths observed by us are using this technique. In this chapter we introduce the general approach to collect statistics of AS path-prepending from used data sets and to find which ASes in which AS-level peerings are prepended. Because the used data sets can not represent the whole Internet, the statistics relying on the data sets are also incomplete.

### 3.1 Approach for Getting AS Path-Prepending

The criterion to evaluate the path-prepending in the Internet is prepending frequency of ASes. Different ASes have different prepending frequency. Some ASes always prepend their local ASNs when announcing to their neighbors due to different reasons like e.g. use of “backup” paths. Those ASes have thus a prepending frequency of 100% while most ASes do not use this path-prepending technique. It is also possible that an AS prepends its ASN by one BGP session while doesn’t do this by another.

The data sets for getting AS path-prepending are BGP table dumps. On the basis of statistics collected from used data sets, we calculate the prepending frequency for each AS-level peering. The prepending frequency for each AS can be easily concluded relying on this result.

### 3.2 Basic Statistics on AS Path-Prepending

An algorithm (Figure 3.2.1) is designed to gain the frequency of prepending. The output of the algorithm is a list of AS-level peerings with prepending frequency in percent.

```
FOR all the records in BGP table dumps
  OBTAIN the AS path from the record
  FOR each ASN X in the AS path
    OBTAIN AS-level peering “Y X”
    INCREMENT Count of prepending for “Y X” by prepending of X
  ENDFOR
ENDFOR

FOR all the AS-level Peerings
  SET Prepending Frequency for the AS-level Peering = Count of Prepending/(Count of
  Prepending + Count of not-prependng)
ENDFOR
```

Figure 3.2.1 Algorithm to get prepending frequency for AS-level peerings

Path-prepending frequency provides insight into the AS path-prepending in the Internet. Table

3.2.1 und Table 3.2.2 presents some basic statistics on the result.

	Number	Percentage
ASes	21,173	100%
ASes sometimes prepended	5,196	24.5%
ASes 100% prepended	378	1.8%
ASes never prepended	15,599	73.7%

Table 3.2.1 Basic statistics on AS path-prepending (concerning ASes)

Table 3.2.1 shows how often ASes are prepended. If an AS prepends its ASN in all AS paths, it is 100% prepended. Those ASes are in the minority (1.8%). 100% prepended AS may happen because this AS doesn't want to be used by the other ASes. Or because our data set is incomplete: the AS is only always prepended in the paths observed by us. The third important reason is use of "backup" AS. In the Figure 3.2.3, AS3 is a "backup" AS of AS4. It prepends its local ASN so that AS path "1 3 3 5 2" is longer than AS path "1 4 5 2". If AS4 falls out, AS3 will take the place of AS4.




Figure 3.2.3 "Backup" AS

*ASes sometimes prepended* are the ASes that are prepended in some AS paths while turn up without being prepended in the other paths. 24.5% of all ASes fall in this group. Beside the both kinds of ASes above, 73.7% ASes are never prepended.

We compute for each AS its prepending frequency and get the density distribution of prepending frequency (Figure 3.2.4). The x-axis is the prepending frequency ranging from 0% to 100% and the y-axis is the density. The diagram illustrates how many ASes are prepended with how much percentage. Most of ASes has a prepending frequency of 0~10%.




Figure 3.2.4 Density of prepending frequency for ASes

ASes are not always prepended for the same times. In the AS path “1 2 2 2 3”, AS2 is prepended two times, in another AS path, AS2 may be prepended one or three or more times. We count all path-prependings in the used data sets and make Table 3.2.2. According to collected statistics, 93.2% ASes are prepended not more than 5 times.

Prepending length	Count of appearance	Percentage	
1	378,139	48,4%	93.2%
2	164,798	21,1%	
3	95,749	12,2%	
4	65,727	8,4%	
5	23,867	3,1%	
>5	53,385	6.8%	
Sum		100.0%	

Table 3.2.2 Statistics on length of AS path-prepending

According to Table 3.2.1 nearly one fourth of ASes are prepended in some AS paths while they are not prepended in the other. So it is a little too coarse if we assume that the ASes are either always prepended or never prepended in the simulation. In order to refine the statistics, we compute prepending frequency for each AS-level peering instead of for each AS:

	Number	Percentage
Peerings	62,595	100%
Peerings sometimes prepended	4,607	7.4%
Peerings 100% prepended	4,792	7.7%
Peerings never prepended	53,196	84.9%

Table 3.2.3 Basic statistics on AS path-prepending (concerning AS-level peerings)

In an AS-level peering like “AS1:AS2”, if AS2 is prepended in all paths that contain this peering, this peering is considered to be always prepended. 7.7% AS-level peerings fall in this group. The reasons for this group are just the same to ASes 100% prepended,

Peerings, by which AS2 is never prepended, are called *peerings never prepended*. Most of the peerings (84.9%) are never prepended. The remaining 7.7% peerings are so-called *peerings sometimes prepended*.

We draw the density distribution of prepending frequency for a better overview (Figure 3.2.5). This diagram indicates how many AS-level peerings are prepended with how much prepending percentage. The x-axis is the prepending frequency ranging from 0% to 100% and the y-axis is the density. Bigger density value indicates that more peerings are prepended with corresponding percentage. One can see that prepending frequency of most AS-level peering is concentrated on 0% and 100%. That is to say, most peerings are either always or never prepended.

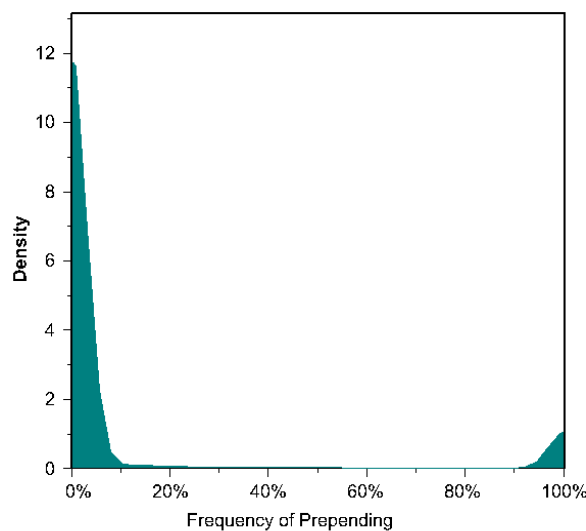


Figure 3.2.5 Density of prepending frequency for AS-level peerings

For the remainder of the work we must know when which ASes should be prepended. The statistics above provide us the necessary information to detect AS path-prepend and configure it accordingly in the simulation.

# 4 Evaluation of AS Path-Prepending Policy through Simulation

The goal of our work is to reproduce BGP routing by means of simulation and to evaluate the role that AS path-prepending policy plays in BGP routing decision. The routing reproduction relies on the AS-level connectivity graph which stems from BGP table dumps. The reproduction must be carried out with as well as without AS path-prepending policy. Both results are compared with the observed AS paths obtained from BGP table dumps in order to disclose the impact of AS path-prepending policy on BGP routing.

This chapter is supposed to explain the approaches and algorithms, which are used to build up the simulation model of the inter-domain routing and to evaluate the influence of AS path-prepending policy. Starting with introduction of the general approach for routing reproduction, the three main steps of simulation are presented in the subsequent parts. In the end a comparison between simulation with and without path-prepending policy is made.

## 4.1 General Approach

The approach to evaluate AS path-prepending policy is divided into 4 main steps (Figure 4.1):

1. We get all the ASes and AS-level peerings from BGP table dumps derived from the Internet. The AS path-prepending policy is applied to the AS-level peerings in which ASNs are always prepended (see section 3.2) such that the simulation is closer to the real Internet. The result of this step is a list of ASes and AS-level peerings that can be considered as an AS-level connectivity graph with policies. This graph consists of nodes (ASes) as well as edges (AS-level peerings) and serves as the input for BGP routing simulation.
2. A simulation model is designed to reproduce the complete BGP routings on the basis of the connectivity graph. This model uses C-BGP as the simulator and the application introduced in section 2.4 to enable large-scale simulation.
3. A comparison between the best AS paths selected by simulation and the observed AS paths in the real Internet is made and the accordance as well as disagreement is recorded.
4. Step 1 to 3 is repeated without AS path-prepending policy. The result is compared with the comparison statistics from Step 3 to disclose the effect that AS path-prepending policy takes on BGP inter-domain routing decision.

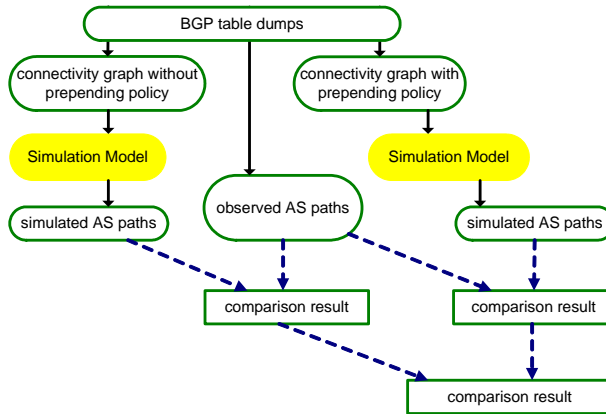


Figure 4.1: General approach for evaluation of AS path-prepending policy

Our research object is the inter-domain routing. An important assumption for our simulation is that each AS has only one router. In the Internet an AS may have many routers whose geographical and functional distribution can take an effect on the BGP routing decision. We neglect this fact because our goal is to find the influence of prepending policy. An approximation of BGP routing is enough to cover the demand.

## 4.2 Construction of AS-Level Connectivity Graph

The configuration files for the simulation refer to AS-level connectivity graph. Those files maintain the list of ASes and AS-level peerings. ASes are the nodes of the connectivity graph and peerings represent edges that connect those nodes.

The first step to construct connectivity graph is to get all the unique ASNs from BGP table dumps and assign each CIDR block to an ASN (See [RFC1519] for more information about CIDR block) so that this AS can originate a prefix during simulation. C-BGP determines the best path from every ASes to the prefixes that are originated by ASNs marked with CIDR blocks in the configuration file. CIDR blocks serve as the address of prefixes. Since each AS should have a unique address we must assign a unique CIDR block to each AS. In our work we use CIDR blocks with a length of 16 bits like “1.0/16”. It is sufficient by reason that ASNs are currently also 16-bit integers [WIKI02]. After the assignment a list of ASNs with CIDR blocks is generated. Those are the nodes in the AS-level connectivity graph.

Now we turn our attention to the AS-level peerings which reveal two kinds of information: between which ASes there exist direct connections and under which policy the BGP sessions take place. We take the AS path “1 2 2 3” as an example, two peerings can be extracted from this path: “1 2” and “2 3” which signify the direct connection between AS1 and AS2 as well as AS2 and AS3. Note that the AS path-prepending is discarded. We read through all AS paths and get a list of AS-level peerings and then apply prepending policy to those peerings.

In the previous capital we get the prepending frequency of every AS-level peering, most of which are 0% or 100%. In order to simplify the simulation, we assume that the peerings are either always (for peerings with a prepending frequency of 100%) or never (for peerings with a prepending frequency of 0%) prepended.

It has to be decided, how many times an ASN should be prepended. We find that ASNs are prepended for different times in BGP table dumps. Typically the prepended ASN is once or twice

prepending. According to the statistics in section 3.2, we set the prepending-times to 5 to maximize the effect of path-prepending, concerning that ISPs bring in AS path-prepending policy in order to make the AS less preferable.

The output of this step is the so-called inter-domain file, which consists of all the peerings with AS path-prepending policy.

The AS connectivity graph constructed in this section is the groundwork of simulation. On the basis of the inter-domain file and domain files, we put BGP routing decision-process into execution with C-BGP.

### 4.3 Implementing Simulation

C-BGP together with the assistant tool described in section 2.4 can simulate BGP routing decision-process for many prefixes. However, during the simulation we find that the number of prefix can affect performance. The time and memory needed for simulation are affected by the number of each time simulated prefixes. Table 4.2 shows the time and memory needed by different number of prefixes for our specific simulation setup:

Number of prefixes simulated each time	CPU time per prefix (approximately)*	Memory (approximately)
1	2 min 30 sec	60MB
3	2 min 20 sec	70MB
4	1 min 44 sec	74MB
5	1 min 37 sec	80MB
6	1 min 40 sec	100MB
10	1 min 50 sec	150MB
50	2 min 10 sec	730MB

\*using 2 AMD Opteron 240 (1.4GHz) simultaneously

Table 4.3 Effect of prefix number on performance

In our work we simulate 5 prefixes from all the more than 20,000 ASes at each time to minimize the simulation time.

We configure the inter-domain policies according to the statistics obtained in section 3.2. For the peerings whose prepending frequency is neither 0% nor 100%, it is difficult to determine their behavior of prepending in the simulation. Fortunately those ASes are in the minority with a share of only 7.7%. In order to simplify our simulation, we consider those ASes as not being prepended while configuring the simulation. Only the peerings 100% prepended in BGP table dumps are configured to be prepended 5 times in the simulation. The other peerings are never prepended.

With the help of the assistant application introduced in section 2.4, we start the simulation with the inter-domain files and one domain file. The whole simulation runs until all the prefixes are simulated, i.e. all domain files are used. Note that the simulation lasts relative long: on our computer (using 2 AMD Opteron 240 (1.4GHz) simultaneously) the whole simulation needs about 2 weeks.

## 4.4 Obtaining Statistics of Simulated Routings

We have got the BGP routing tables in the last section. In this section, we compare AS paths from simulations both with and without policy respectively with the observed paths in the Internet. Both comparison results are again compared to disclose the impact of AS path-prepending policy on BGP routing decision.

In the simulated routing tables, for each connection from one AS (e.g. AS1) to another (e.g. AS2) there may be two kinds of paths (Figure 4.4.1):

- The best path from AS1 to AS2
- Alternative path(s) from AS1 to AS2

The best path is the path that is actually used in the simulated topology. Alternative paths are the paths that are used only if the best path fails.

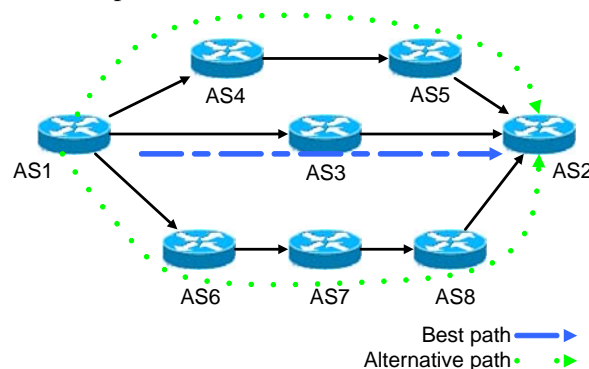


Figure 4.4.1: Categories of simulated AS paths

Each time we take a path (referred to as observed path) from the BGP table dumps and search for the same path in the simulated routing tables. There are 3 possible results:

- Case 1: The observed path is in accordance with the simulated best path. (Figure 4.4.2)
- Case 2: The observed path is in accordance with a simulated alternative path instead of the best path (Figure 4.4.2)
- Case 3: Neither the best path nor a alternative path is the same to the observed one.

Note that the comparison discards the path-prependings both in BGP table dumps and simulated result. I.e. path “1 3 2” and path “1 3 3 2” are considered to be the same path.




Figure 4.4.2 Possible results of paths comparison

The number of Case 1, 2, 3 is recorded during the comparison. Table 4.4.1 shows the sums of them.

	simulation without AS path-prepend policy		simulation with AS path-prepend policy	
Best path match	1,220,183 (25.8%)	2,561,389 (54.2%)	1,199,396 (25.4%)	2,513,138 (53.2%)
Alternative path match	1,341,206 (28.4%)		1,313,742 (27.8%)	
No match	2,163,333 (45.8%)		2,211,584 (46.8%)	
Sum	4,724,722 (100%)			

Table 4.4.1 Comparison of simulation result

The simulation result is very different from the real Internet. There are a number of reasons. For example, our data set does not cover the whole internet and we don't consider intra-domain routing. However, one reason is the most important. BGP routes are associated with many attributes. These BGP attributes are considered during the BGP decision process when a best route is selected from the set of all routes learned for a certain prefix. In the Internet, it is possible to define a set of filter rules for each BGP session. A router may manipulate some attributes of the route according to the filter rules and thus influence BGP routing decisions. In our simulation we consider only *AS path* and the other attributes like for example Local Preference and *Multiple Exit Discriminator (MED)* are not used.

As we can see in Table 4.4.1, if we don't bring in AS path-prepend policy, 54.2% observed paths turn up in the simulated result. Under the AS path-prepend policy, the percentage

descends to 53.2% on the contrary.

The simulation under AS path-prepend policy generates fewer best paths that exist also in BGP table dumps than the simulation without the policy (25.4%:25.8%). Although the difference is relative small, this result is different from our anticipation. We expect that AS path-prepend policy can make the result closer to the real Internet.

In the next subsection we try to explain the reasons for this result.

## 4.5 Analysis of Simulation Result

We bring path-prepend policy into our work, expecting that it can make the simulation closer to the real Internet. However, the result shows that the simulation without policy is even better than that with policy. In order to find the reason for this result, we try to explain the following 4 phenomena in the results:

1. Some observed paths are in accordance with the best paths generated by simulation with prepending policy while they exist only as alternative paths by simulation without prepending policy. Why?

We take an example to explain this phenomenon. AS path “1 3 4 2” is the route from AS1 to AS2 in BGP table dumps. In the result of simulation without the prepending policy, this path only serves as an alternative path. The best one is “1 5 2”. However, in the simulation under prepending policy, AS path “1 3 4 2” takes the place of the best path. The reason is that in the AS-level connectivity graph, which serves as the basis of simulation, there is a peering “1 5” by which AS5 is prepended 5 times if it uses prepending policy. Therefore, the AS path “1 5 2” changes into “1 5 5 5 5 5 2” which is longer than AS path “1 3 4 2”, so the later is selected. Figure 4.5.1 shows the decision-process.




Figure 4.5.1 Phenomenon 1: Policy converts alternative path into best path

2. Reciprocally, some observed paths are in accordance with the best paths generated by

simulation without prepending policy while they exist only as alternative paths by simulation with prepending policy. Why?

Once again we take an example. In BGP table dumps AS path “1 3 3 2” is the best route from AS1 to AS2 in spite of the path-prepending by AS3. In our simulation, AS3 is always prepended 5 times in the peering “1 3” and the path under prepending policy is transformed into “1 3 3 3 3 3 2” so that AS path “1 4 5 2” turns out to be better and is however not in accordance with the observed path “1 3 3 2” any more. On the contrary, if we don’t use prepending policy, the traffic is sent across the same ASes as the observed path.

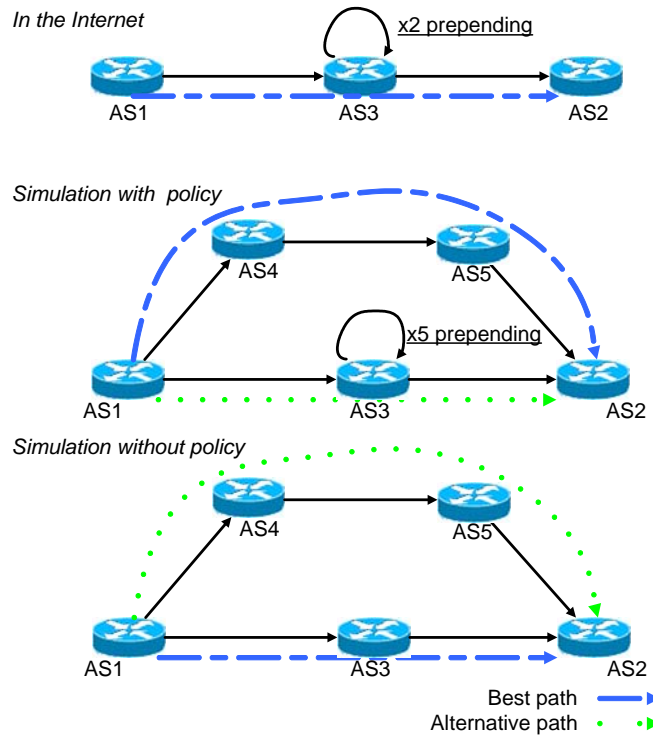


Figure 4.5.2 Phenomenon 2: Policy converts best path into alternative path

3. Some observed paths exist in the routing tables generated by simulation with prepending policy while they don’t turn up by simulation without prepending policy. Why?

The reason is as follows. Take the path “1 3 5 2” in the Internet for an example. In the routing tables of simulation under the prepending policy, this path serves as either the best path or an alternative path. Although “1 3 4 4 4 4 4 2” is also a possible route, AS3 hasn’t learnt the route from itself to AS4 because the simulator finds “3 5 2” shorter than “3 4 4 4 4 4 2”. By the simulation without policy the simulator selects “1 3 4 2” since “3 4 2” is considered better than “3 5 2” (there may be different reasons, e.g. because ASN 4 is smaller than ASN 5) and the AS path “1 3 5 2” is not learnt. Figure 4.5.3 shows the decision-process.




Figure 4.5.3 Phenomenon 3: Policy converts path not learnt into best path

4. Some observed paths exist in the routing tables generated by simulation without prepending policy while they don't turn up by simulation with prepending policy. Why? The reason is as follows. The AS path "1 3 4 4 2" is the route from AS1 to AS2 in the Internet in spite of the path-prepending by AS4. In the routing tables of simulation under the prepending policy, this path changes into "1 3 4 4 4 4 4 2" and "1 3 5 6 2" turns out to be shorter so that the route from AS3 to AS4 is not learnt. By the simulation without policy the simulator selects "1 3 4 2" as a possible path. Figure 4.5.4 shows the decision-process.




Figure 4.5.4 Phenomenon 4: Policy converts best path into path not learnt

The 4 phenomena referred to above are the reasons for the different results of simulations with and without path-prepending policy.




Figure 4.5.5 Impact of policy on simulated AS paths

Figure 4.5.5 shows how exactly the impact of path-prepending policy on simulated BGP routing is. We see in this diagram that 19,198 best AS-Paths are degraded to be alternative ones and 78,707 can't be learnt any more because of use of prepending policy. On the other hand, 36,332 alternative paths and 41,478 paths not learnt are selected as the best paths as a result of use of prepending policy. That means the sum of best paths decreases by using policy since  $36,332 + 41,478 = 77,810$  is less than  $19,198 + 78,707 = 98,625$ . As a result, simulation with policy generates more AS paths that are in accordance with the observed ones.

The other data difference in Table 4.4.1 can also be explained in the same way with the help of Figure 4.5.5. It is a result of interaction of the 4 phenomena. Phenomena 2 and 4 are the reasons why simulation result without policy is better than that with policy. The observed paths in the Internet are selected as the best path in spite of path-prependings in them. There are different reasons. E.g. there exist other policies that we don't know and can't consider. In the simulation with path-prepending policy, the best paths are replaced by the other paths by maximizing the effect of path-prependings.

## 5 Conclusion and Future Work

The goal of our work consists of obtaining statistics of path-prependings in the Internet and evaluating the impact of prepending policy on routing decisions by means of simulation. In this work, we have described our approach to build a realistic model of the inter-domain routing. As ISPs generally do not disclose topology information, paths gathered by BGP route collectors are used to obtain ASes and peerings between ASes. Relying on this AS-level connectivity graph, we start simulation of BGP routing decisions with and without prepending policy. The generated routing tables of both simulations are compared to reveal the impact of prepending policy on BGP routing decisions.

The results are surprising. Although path-prependings policy is widely used in the Internet, its influence can't be seen in the results. New generated routes by simulation without path-prependings policy are closer to the real Internet than those simulated with path-prependings policy. The analysis in Subsection 4.4 and case study in Subsection 4.5 shows some reasons for this result. The basic reason is that BGP routing decision is influenced by many factors. Our simulation doesn't consider all the factors. It is incomplete. E.g. we don't consider other BGP policies or use of "backup" paths. There are also other reasons, like the used data sets can't represent the total Internet and we don't simulate intra-domain routing which can also play a role in routing decisions.

We believe that if we refine the simulation by collecting more raw data sets from the Internet, by considering more BGP policies as well as by finding other possible factors that can affect BGP routing decision, the simulation result with policy will be closer to the Internet than that without policy and the impact of path-prependings policy on routing decisions can be seen.

## 6 List of Tables

Table 2.1.1 An example of AS path.....	7
Table 2.2 Basic statistics on used raw data (Excluding private ASes) .....	8
Table 3.2.1 Basic statistics on AS path-prepend (concerning ASes).....	11
Table 3.2.2 Statistics on length of AS path-prepend .....	12
Table 3.2.3 Basic statistics on AS path-prepend (concerning AS-level peerings).....	12
Table 4.3 Effect of prefix number on performance .....	16
Table 4.4.1 Comparison of simulation result .....	18

## 7 List of Figures

Figure 2.1 Internet routing protocols .....	6
Figure 2.1.2 Effect of AS path-prependng on BGP routing selection .....	8
Figure 3.2.1 Algorithm to get prepending frequency for AS-level peerings .....	10
Figure 3.2.3“Backup” AS.....	11
Figure 3.2.4 Density of prepending frequency for ASes.....	12
Figure 3.2.5 Density of prepending frequency for AS-level peerings.....	13
Figure 4.1: General approach for evaluation of AS path-prependng policy .....	15
Figure 4.4.1: Categories of simulated AS paths .....	17
Figure 4.4.2 Possible results of paths comparison .....	18
Figure 4.5.1 Phenomenon 1: Policy converts alternative path into best path.....	19
Figure 4.5.2 Phenomenon 2: Policy converts best path into alternative path.....	20
Figure 4.5.3 Phenomenon 3: Policy converts path not learnt into best path .....	21
Figure 4.5.4 Phenomenon 4: Policy converts best path into path not learnt .....	22
Figure 4.5.5 Impact of policy on simulated AS paths .....	22

## 8 Bibliography

- [CR01] M. Caesar AND J. Rexford, BGP routing policies in ISP networks, available at: <http://www.eecs.berkeley.edu/Pubs/TechRpts/2005/6507.html>, 2005.
- [FE01] A. Feldmann et al., Netscope: Traffic Engineering for IP Networks, IEEE, Network, March 2000.
- [FWR01] N. Feamster, J. Winick, and J. Rexford, A Model of BGP Routing for Network Engineering, Proc. ACM SIGMETRICS, June 2004.
- [MU01] W. Mühlbauer, Towards Reproducing Inter-Domain AS Paths, Master Thesis, November 2005.
- [PA01] J. Papen, Demestifying BGP, Linux Magazine, May 2003.
- [QT01] B. Quoitin AND S. Tandel, C-BGP User's Guide, Université Catholique de Louvain Louvain-la-Neuve, Belgium, July 12, 2005.
- [QU01] B. Quoitin, C-BGP - An efficient BGP simulator (version 1.20), available at: <http://cbgp.info.ucl.ac.be>, July 12, 2006.
- [QU02] B. Quoitin AND S. Uhlig, Modeling the routing of an Autonomous System with C-BGP, IEEE Network, Vol 19(6), November 2005.
- [RFC1519] V. Fuller, T. Li, J. Yu AND K. Varadhan, Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy, RFC 1519, September 1993.
- [RFC4271] Y. Rekhter, T. LI AND S. Hares, A Border Gateway Protocol 4 (BGP-4), RFC4271, January 2006.
- [WIKI01] Wikipedia Contributors, Internet backbone, Wikipedia, The Free Encyclopedia, available at: [http://en.wikipedia.org/wiki/Internet\\_backbone](http://en.wikipedia.org/wiki/Internet_backbone), July 14, 2006.
- [WIKI02] Wikipedia Contributors, Autonomous System (Internet), Wikipedia, The Free Encyclopedia, available at: [http://en.wikipedia.org/w/index.php?title=Autonomous\\_system\\_%28Internet%29&oldid=59756304](http://en.wikipedia.org/w/index.php?title=Autonomous_system_%28Internet%29&oldid=59756304), July 8, 2006.
- [WIKI03] Wikipedia Contributors, Border Gateway Protocol, Wikipedia, The Free Encyclopedia, available at: <http://en.wikipedia.org/wiki/Bgp>, July 8, 2006.