

# Understanding Current IPv6 Performance: A Case Study from CERNET

Yi Wang<sup>1</sup>, Shaozhi Ye<sup>2</sup>, Xing Li<sup>3</sup>

Department of Electronic Engineering, Tsinghua University,  
Beijing 100084, P. R. China

<sup>1</sup>wangyi@ns.gtest.edu.cn

<sup>2</sup>ys@compass.net.edu.cn

<sup>3</sup>xing@cernet.edu.cn

**Abstract.** Much work has been done on IPv6 standards and many IPv6 testbeds have been deployed. However, little is known about the performance of the real IPv6 Internet, especially from the perspective of end users. In this paper, we perform a case study of current IPv6 performance from CERNET<sup>1</sup>. We propose an active measurement methodology, which analyzes the traces of file transference from IPv4/IPv6 dual-stack Web servers. We study 936 dual-stack Web servers located in 44 countries by collecting and analyzing over 585,680 IPv4/IPv6 packet-level traces with 133,340 million packets. A comprehensive performance comparison of IPv4 and IPv6 is presented, including connectivity, packet loss rate, round-trip time and etc. We also present a brief case study on relationship of RTT and network topology, which is helpful to improve the performance of IPv6 networks. Finally, we discuss the generality and speciality of our CERNET case and results. To our best knowledge, this paper is the first performance study based on large scale IP traffic measurement in real IPv6 Internet.

## 1 Introduction

IPv6 is proposed by IETF to provide the Internet with larger address space and better performance [1]. In the past ten years, a lot of work has been done on the protocol design [4], connection and routing mechanism [5], [6], [7], and transition mechanisms [8], [9] of IPv6. As the demand of IPv6-supported network equipments increases, some performance evaluation methods and platforms are proposed, which mainly focus on the performance of hardware and its compatibility with IPv6 protocols [10], [11].

After its first decade of protocol design and testing, the IPv6 Internet is now in a transition phase from experimental research networks to global operational networks. Once a network begins to provide public services, its performance is always a big issue. As to IPv6, it is an even more complex question mark. On one hand, IPv6 Internet is not constructed on a clear map. It faces transition problem from the very

---

<sup>1</sup> China Education and Research Network

beginning. Tunneling plays an important role in IPv6 testbeds and the migration from IPv4 to IPv6 and IPv6-in-IPv4 tunnels are widely used where no native IPv6 connectivity available. However, it is now believed that tunnels degrade the network performance and reliability [2], [3]. On the other hand, little is known about the actual performance of the growing IPv6 Internet, especially from the perspective of an average IPv6 end user. We believe that end users are responsible for a large fraction of Internet traffic, and their experiences of network performance are helpful to the design and deployment of IPv6 networks.

Recently, more attention is paid to the performance and operational issues of IPv6 networks [2], [3]. In [2], the authors discuss the IPv6-in-IPv4 tunnel discovery issue, and propose a set of techniques to infer tunnels. Each technique is a combination of basic methods: Path MTU discovery, DNS lookups, IP spoofing, hop limit manipulation and IPv6 header modifying. Their experimental results show that even “native” networks reach more than 60% of all IPv6 prefixes through tunnels. The authors of [3] argue that poorly managed experimental IPv6 sites are one of the major hurdles to the perceived quality of the IPv6 Internet. With focuses on troubleshooting, they select a group of IPv4/IPv6 dual-stack nodes by DNS lookups. They study the IPv6:IPv4 RTT ratios using dual-stack *ping* and do path analysis using *traceroute* with Path MTU discovery from three different locations in Japan and Spain.

The novelty of our work in this paper is emphasized by the fact that no previous work attempts to characterize the performance of IPv6 Internet by real TCP traffic measurement. Besides, by studying the relationship between RTT and network topology, we briefly propose an effect method to locate the key point of IPv6 performance for configuration and troubleshooting.

The reminder of this paper is organized as follows. Sect. 2 proposes the methodology of IPv4/IPv6 dual-stack Web server measurement. In Sect. 3 we present the measurement results as well as our analysis of the IPv6 Internet performance. We discuss the generality and speciality of our CERNET case and some possible factors that may affect the results in Sect. 4. Sect. 5 concludes the whole paper by summarizing our contributions and providing directions for the future work.

## 2 Methodology

As an active measurement study, we design our methodology as follows. First, we obtain a list of dual-stack Web servers from our IPv6 Web search database which has been collecting information of IPv6 Web sites since 2001. Then we perform large scale data gathering by crawling web pages and logging the connection traces. We also run *ping* and *ping6* tests to study the performance differences between IP packets and ICMP packets.

### 2.1 Dual-Stack Web Server List

WWW (more precisely, the HTTP on port 80) is one of the most widely used service in current IPv4-based Internet [12]. For the purpose of smooth transition from IPv4

to IPv6, more and more Web servers are implemented and configured with both IPv4 and IPv6 protocol stacks. We choose these dual-stack Web sites as our data sources since we can gain better understanding of IPv6 performance as well as its distinctive problems by comparison with its IPv4 counterpart.

We have been monitoring the evolution of Web sites on the IPv6 Internet since May 2001 and found that the number of IPv6 accessible Web sites keeps growing steadily [13]. For the experiment in this paper, we use an initial hostname list of 1,306 IPv6 Web sites. These Web sites has been observed on service at least for one month since May 2001. On Sept.1<sup>st</sup> 2004, the beginning of our experiment, 1,235 of them can still be resolved to either IPv6 or IPv4 addresses (or both). In this paper, we are interested in those dual-stack Web servers with both IPv6 and IPv4 global routable addresses. After removing duplicates with identical IPv6 and IPv4 address pairs but different hostnames, we finally obtain a list of 936 dual-stack Web servers, which is used for our data collection.

## 2.2 Design of Testbed and Experiment

In order to imitate the experience of average IPv6 end users, our testbed consists of a PC running FreeBSD 4.10 operating system with IPv4 and IPv6 dual stacks, as well as 100Mbps Ethernet links to both IPv4 and native IPv6 Internet through CERNET.

During the experiment, we run two programs, *wget* [14] and *tcpdump* [15]. We use *wget* and its IPv6 patched version *wget6* to download the same files from Web servers through IPv4 and IPv6 links respectively. The first level files of Web sites are often the *index.htm* pages, which are usually no more than several Kbytes. To avoid being biased by short connections, we crawl two or three levels deeper for larger files. We use a script to start 20 processes of *wget* and *wget6* together and make the IPv4 and IPv6 crawling jobs of the same sites working at almost the same time. Meanwhile, we make two *tcpdump* listen to the IPv4 and IPv6 interfaces respectively and log traces. We use a *cronjob* to collect data and dump traces periodically.

## 3 Experimental Results

### 3.1 Experiment Summary


From Sept.1<sup>st</sup> to Sept.7<sup>th</sup>, 2004, we performed Web crawling and trace logging once a day with large amount of data transference. From Sept.8<sup>th</sup> to Sept.10<sup>th</sup>, we did six consecutive measurements everyday, with less data collected each time. Altogether, we collected over 585,680 connection traces of about 133.34 million packets from the 936 dual-stack web servers. In terms of bytes, about 63.8G bytes data are from IPv6 connections and the other 109.2G bytes data are from their IPv4 counterparts.

Table 1 shows the distribution of the dual-stack nodes by their country codes, representing 44 countries. (It is possible that the real location of a node is different from

the registered country.). We obtain the country codes for 936 nodes from *whois* databases of APNIC, RIPE and ARIN.

**Table 1.** Distribution of dual-stack web servers by country codes

JP: 298	GB: 32	FI: 16	ES: 10	TW: 8	KR: 6	MX: 3	PH: 2	IN: 1
NL:133	IE: 28	CA: 13	MY: 10	ZA: 8	BE: 5	TH: 3	RU: 2	LU: 1
US: 83	FR: 20	DK: 13	PL: 10	CN: 7	GR: 5	BR: 2	BG: 1	UK : 1
DE: 77	CH: 16	AU: 12	NO:9	HU: 7	PT: 4	ID: 2	HK: 1	YU: 1
IT: 39	EE: 16	CZ: 12	SE: 8	AT: 6	SK: 4	IS: 2	HR: 1	



**Fig. 1.** Distribution of hop counts of IPv4 and IPv6 connections in the experiment

Fig. 1 shows the hop counts distribution of the connections in the experiment. In our experiment, the average hop count is 8.7 in IPv6 Internet, and 17.5 in IPv4 Internet. Moreover, the hop count in IPv4 tends to be diversified in a wider range than in IPv6. The simplicity of topology and the existence of tunnels may account for these discrepancies.

### 3.2 Connectivity

Connectivity is a fundamental requirement to a network. In the current IPv4 Internet, most of the time connectivity is not such a big problem as it was decades ago. However, it might be an issue in the IPv6 Internet.

To examine the connectivity of the 936 dual-stack Web servers, we use both the traditional *ping/ping6* and *wget* to test their connectivity. There is possibility that some server is up and responses to ICMP *ping* packets, but has nothing on service at

port 80, or vice versa (probably the *ping* service is disabled). Table 2 compares the changes of connectivity of the 936 servers in a day by *ping/ping6*. It shows that the accessibility of IPv6 Web servers tend to be more variable during the day.

**Table 2.** Average reachability of dual-stack Web servers by *ping/ping6* in a day

	IPv4 Accessible	IPv6 Accessible	IPv4 & IPv6 Accessible	IPv4 & IPv6 Inaccessible
0:00 – 6:00	774 (82.7%)	728 (77.8%)	639 (68.3%)	73 (7.8%)
6:00 – 12:00	772 (82.5%)	718 (76.7%)	629 (67.2%)	75 (8.0%)
12:00 – 18:00	771 (82.4%)	697 (74.5%)	609 (65.1%)	77 (8.2%)
18:00 – 24:00	772 (82.5%)	626 (66.9%)	636 (67.9%)	75 (8.0%)

**Table 3.** Average accessibility of dual-stack Web servers by *ping/ping6* in different regions

IPv6 IPv4		Accessible Accessible	Accessible Inaccessible	Inaccessible Accessible	Inaccessible Inaccessible
Total	936 (100%)	642 (68.6%)	89 (9.5%)	132 (14.1%)	73 (7.8%)
RIPE	478 (100%)	312 (65.3%)	60 (12.5%)	76 (15.9%)	30 (6.3%)
JP	298 (100%)	235 (78.9%)	23 (7.7%)	24 (8.1%)	16 (5.3%)
ARIN	108 (100%)	68 (63.0%)	2 (1.8%)	24 (22.2%)	14 (13.0%)
APNIC	52 (100%)	27 (41.9%)	4 (7.7%)	8 (15.4%)	13 (25.0%)

**Table 4.** Average accessibility of dual-stack Web servers by *wget/wget6* in different regions

IPv6 IPv4		Accessible Accessible	Accessible Inaccessible	Inaccessible Accessible	Inaccessible Inaccessible
Total	936 (100%)	584 (62.4%)	73 (7.8%)	191 (20.4%)	88 (9.4%)
RIPE	478 (100%)	289 (60.4%)	53 (11.1%)	94 (19.7%)	42 (8.8%)
JP	298 (100%)	202 (67.8%)	17 (5.7%)	52 (17.4%)	27 (9.1%)
ARIN	108 (100%)	63 (58.3%)	2 (1.9%)	30 (27.8%)	13 (12.0%)
APNIC	52 (100%)	30 (57.7%)	1 (1.9%)	15 (28.8%)	6 (11.5%)

Table 3 summarizes the accessibility of the 936 servers by partitioning them into four general regions. Table 4 shows the results of *wget/wget6* experiment. Since Japan contributes a large fraction of servers with relatively close locations, we separate them from the non-JP APNIC servers as [3] does. We also merge the small number of LACNIC servers into the ARIN ones. The Japanese servers perform best connectivity, since majority IPv6 links from CERNET to RIPE, ARIN and other APNIC nodes first go through Japan. In Table 3, APNIC servers have exceptional low connectivity by *ping* for both IPv4 and IPv6. Comparing with Table 4, it is probably because many APNIC Web servers have disabled the *ping* service.

### 3.3 Packet Loss

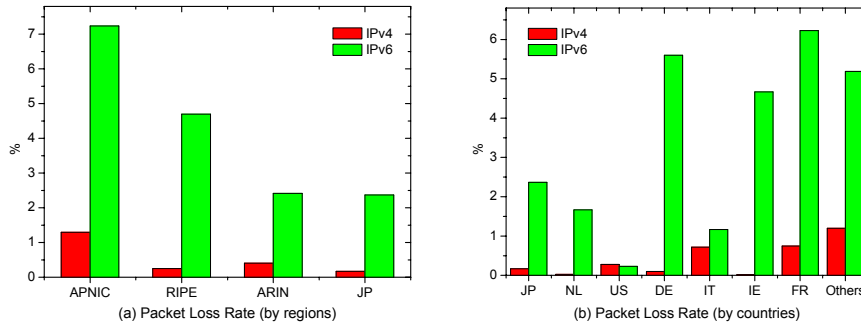
Due to the development and enormous diversity of the Internet, average packet loss rate in different studies is reported in a large range. It is between 0.36% and 3.54% by Boralla et al. [16] based on the study of speech data transmission, between 1.38% and

11% Yajnik et al. [17] based on the measurement at the Mbone receiver, between 2.7% and 5.2% by Paxson [18] based on his long-term experiment of bulk data transference. The most similar data source as our paper may be the one used by Balakrishman et al. [19]. They analyzed the dynamics of a large number of TCP web sessions at a busy Web server, and reported an average packet loss rate of 0.49%.

In our experiment, the IPv6 and the IPv4 connections have an average packet loss rate of 3.09% and 0.76% respectively. Table 5 shows the detailed distribution. We notice that the IPv6 links tend to introduce more packet loss than the IPv4 ones as a whole.

**Table 5.** Packet loss distribution by *wget/wget6* measurement

IPv4	No loss occurs	No loss occurs	Loss occurs	Loss occurs
IPv6	No loss occurs	Loss occurs	No loss occurs	Loss occurs
Percent	49.3%	27.8%	8.8%	14.1%



**Fig. 2.** Distribution of packet loss rate

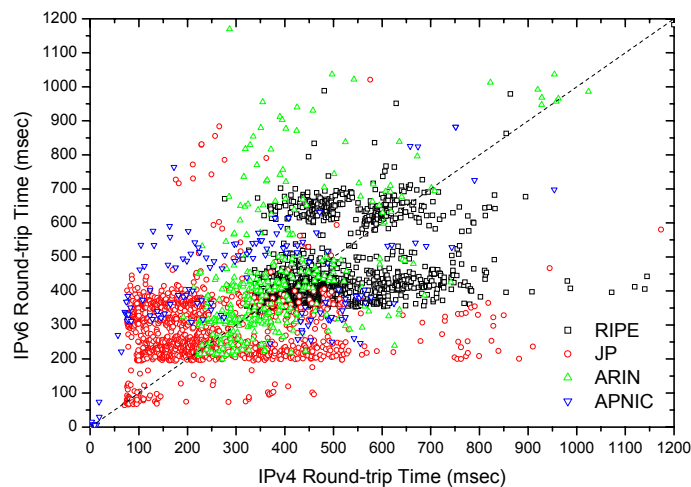
Fig. 2 shows the packet loss rate of the four general regions and some representative countries. We notice that the IPv6 packet loss rates of DE (Germany), IE (Ireland) and FR (France) are extraordinarily high. After examining the corresponding traces, we find that there are one or more “special links” from each of these countries to our testbed. These links are of large volume of data but with quite high packet loss rate (exceed 10%, sometimes even 20%), so they increase the average loss rate of the country as a whole. However, some other links using the same BGP routing table do not experience such high loss rate, so we argue problems of these “special links” probably occur in the access links near the Web servers, not in the IPv6 backbone.

### 3.4 Round-trip Time

Round-trip time (RTT) is another important parameter to indicate the quality-of-service (QoS) of a network. Moreover, in the scenario of this paper we find its distri-

bution also an insightful tool to reveal the network topology and deployment information.

Fig. 3 shows the distribution of the observed RTTs in scatter plot. It plots about 3,600 RTT value pairs of the Web servers accessible by both IPv4 and IPv6 web crawlers in the 10-day measurements. For each value pair, the IPv4 RTT is plotted on the X-axis and the IPv6 RTT is plotted on the Y-axis. We also draw the unity line,  $y = x$  to help the comparison.



**Fig. 3.** Scatter plot of IPv4/IPv6 RTT distribution

We find in Fig. 3 that different regions tend to have different typical RTT value ranges, which accords with our intuition. The RIPE nodes have the most concentrated RTT distribution, considering its largest total number. Interestingly, they cluster into two groups with IPv6 RTT range approximately from 350 ms to 450 ms and from 600 ms to 700 ms respectively. The Japanese nodes also exhibit this characteristic, which indicates the simplicity of the current IPv6 Internet topology. Many Web servers routed differently in IPv4 Internet are visited through the same IPv6 backbone. Moreover, the two clusters with different IPv6 RTT ranges are probably the results of native IPv6 connection and IPv6-in-IPv4 tunneling. The ARIN nodes do not have such a notable clustering characteristic as the RIPE and Japanese nodes. The majority of the ARIN nodes are around the unity line. In spite of their small total number, the APNIC nodes have large variance of RTT values due to their topology diversity.

Fig. 4 shows the probability distribution function (PDF) of the IPv4/IPv6 RTTs. It also reveals that the distribution of IPv6 RTTs tends to be more concentrated than the IPv4 RTTs. The three major peaks of the IPv6 PDF curve correspond to the three major IPv6 RTT clusters in Fig. 3.




Fig. 4. Distribution of RTT of IPv4 and IPv6 connections

### 3.5 Case Study

To further explore the cause of “clustering” of IPv6 RTTs, we study the relationship between IPv4/IPv6 RTT values and AS topology. We partition the IPv6 addresses by matching them with the prefixes of CERNET IPv6 BGP table [20]. There were 658 prefixes in the table when processing our data. Combining the RTTs with the network topology provides valuable information for both network configuration and troubleshooting. Limited by the space of this paper, we just show an example of identifying the “hinge” of IPv6 performance in the AS map.

When analyzing the traces, we find that all the IPv6 addresses going through the AS path: “4538-X-X-X-X-2914-X-i” have RTT values about 320 ms to 420 ms. However, their IPv4 counterparts have a very large RTT value range, from 230 ms to 600 ms. This comparison yields insight into the path differences and suggests that these IPv6 paths do not follow their IPv4 counterparts. We verify this guess by using the *traceroute* and *traceroute6*. To further identify the property of the IPv6 paths (e.g., whether they are native links or tunneled links) may require more detailed information such as Path MTU, which is out of the scope of this paper and is mentioned as the future work in Sect. 5.

## 4 Discussion

Fig. 5 shows the current CERNET IPv4 and IPv6 international connectivity. The majority IPv6 connections from CERNET to RIPE, ARIN and APNIC first go through the APAN-JP link to Japan, so the performance of this link is critical to our

measurement. Nevertheless, from Fig. 2 and Fig. 3 we can see the IPv6 performance (in terms of packet loss and RTT) of ARIN nodes is distinct from JP nodes and RIPE nodes, even sharing the same APAN-JP link. So we suggest that different routing policies of IPv6/IPv4 in CERNET do not affect our measurement results notably.

Besides, the performance difference between IPv6 servers and IPv4 servers of the same site, and the different processing speeds to IPv6 and IPv4 packets of routers both may introduce biases. However, we intend to test neither the high-load performance of the IPv6/IPv4 web servers nor the routers themselves, and the data rates in our measurement are as normal as average end users surfing the web. As a result, the influence of router and web server performances to our results should be very small.

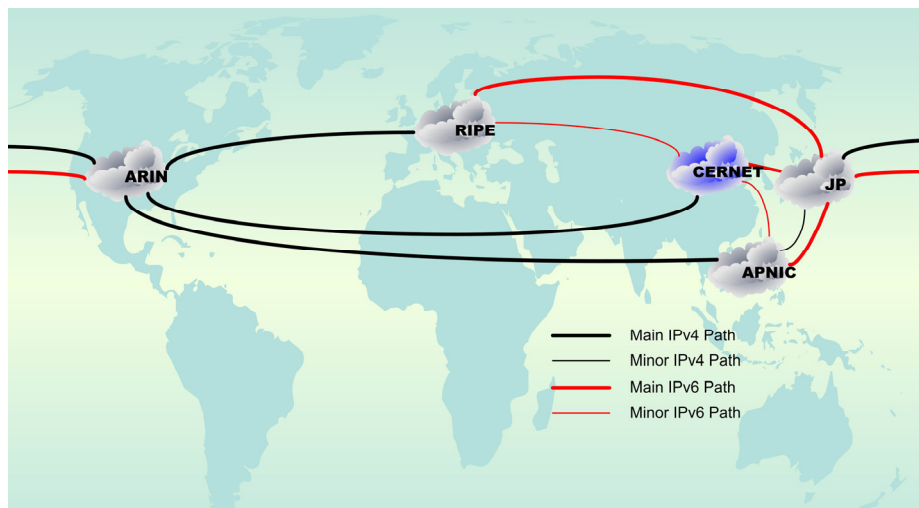


Fig. 5. CERNET IPv4 and IPv6 international connectivity

## 5 Conclusions

Today, IPv6 is in an important transition state from testbeds to operational networks. Future IPv6 development and engineering require clear understanding of current IPv6 performance. In this paper, we design an active measurement methodology to compare IPv6 and IPv4 performance from the perspective of a common end user. Our methodology includes obtaining and screening of IPv6/IPv4 dual-stack Web servers as well as Web file crawling and trace logging techniques. We perform our measurement from native IPv6 network in CERNET to 936 dual-stack web sites worldwide and present the first experimental results of IPv6 performance based on IP traffic in real Internet. Our analytical results indicate that there is still space left to improve the IPv6 performance, which includes increasing its connectivity and reducing its packet

loss rate. We also consider the distribution of IPv6/IPv4 RTT value pairs a good indicator to locate the key point of performance with the help of path information.

Our effort to understand the IPv6 performance and operational issues is not completed. We are interested in the performance differences of native and tunnel links of IPv6. A systematic comparative study using tunnel discovery techniques is on-going. We also plan to follow its performance evolution in larger scope, and study its interaction with IPv4 Internet by measurement and modeling.

## References

1. S. Bradner and A. Mankin: IP: Next Generation (IPng) White Paper Solicitation. RFC 1550, Dec. 1993
2. L. Colitti, G. D. Battista, and M. Patrignani: IPv6-in-IPv4 tunnel discovery: methods and experimental results. *IEEE Transactions on Network and Service Management*, vol. 1, no.1. April. 2004
3. K. Cho, M. Luckie, B. Huffaker: Identifying IPv6 Network Problems in the Dual-Stack World. SIGCOMM '04 Network Troubleshooting Workshop, Portland, USA, 2004
4. S. Deering and R. Hinden: Internet Protocol, version 6 (IPv6) Specification. RFC 2460, Dec. 1998
5. R. Rockell and R. Fink: 6Bone Backbone Routing Guidelines. RFC2772, Feb. 2000
6. B. Carpenter and K. Moore: Connection of IPv6 Domains via IPv4 Clouds. RFC 3056, Feb. 2001
7. A. Durand, P. Fasano, I. Guardini, and D. Lento: IPv6 Tunnel Broker. RFC 3053, Jan. 2001
8. R. Gilligan and E. Nordmark: Transition Mechanisms for IPv6 Hosts and Routers. RFC 2893, Aug. 2000
9. Ioan Raicu and Sherali Zeadally: Evaluating IPv4 to IPv6 Transition Mechanisms, *IEEE International Conference on Telecommunications 2003, ICT'2003, Tahiti Papeete, French Polynesia*, Feb. 2003
10. Internet Protocol version 6 (IPv6) Conformance and Performance Testing. [http://www.ixiacom.com/library/white\\_papers/wp\\_display.php?skey=ipv6](http://www.ixiacom.com/library/white_papers/wp_display.php?skey=ipv6)
11. IPv6 Testing. <http://www.spirentfederal.com/technology.asp?MS=67>
12. S. McCreary and K. Claffy: Trends in Wide Area IP Traffic Patterns: A View from Ames Internet Exchange. May 2000. <http://www.caida.org/outreach/papers/2000/AIX0005/>
13. Yue Li, Hui Liu, Gang Zhu, Shaozhi Ye, and Xing Li. Analysis of IPv6 over Search Engine. The Fifth Joint AEARU Workshop on Web Technology and Computer Science, Oct 2003
14. GNU wget. <http://www.gnu.org/software/wget/wget.html>
15. TCPDUMP. <http://www.tcpdump.org/>
16. M.S. Borella, D. Swider, S. Uludag, and G.B. Brewster: Internet Packet Loss: Measurement and Implications for End-to-End QoS. *International Conference on Parallel Processing*, Aug. 1998
17. M. Yajnik, S. Moon, J. Kurose, and D. Townsley: Measurement and Modeling of the Temporal Dependence in Packet Loss. *IEEE INFOCOM*, March 1999
18. V. Paxson: Measurements and Analysis of End-to-End Internet Dynamics. Ph.D. dissertation. University of California at Berkeley, 1997
19. H. Balakrishnan, V.N. Padmanablah, S. Seshan, M. Stemm, and R.H. Katz: TCP Behavior of a Busy Internet Server: Analysis and Improvements," *IEEE INFOCOM*, March 1998
20. CERNET BGP VIEW Project. <http://bgpview.6test.edu.cn/bgp-view/index.shtml>