Characteristics of Internet Traffic in Thailand

Panita Pongpaibool

NECTEC 112 Phahol Yothin Rd., Klong Luang, Pathumthani 12120 THAILAND, panita@nectec.or.th

ABSTRACT

Understanding Internet traffic behaviors is essential to planning and management of existing networks, as well as to designing next generation networks. While there are many publications on patterns and characteristics of Internet traffic elsewhere in the world, little has been published about behaviors of Internet traffic in Thailand. This article reports on measurements from three major Internet Exchanges in Thailand where all Internet service providers are required to connect to. We reveal the characteristics of traffic at these locations in terms of short-term behaviors and long-term growth trends. We also discuss the impact of our findings to the Thai Internet industry.

Keywords: Internet traffic, Internet growth, traffic pattern, traffic measurement

1. INTRODUCTION

Understanding Internet traffic behaviors is crucial for network planning, and for understanding the health of the telecom industry. While there are many publications on patterns and characteristics of Internet traffic in other countries, little has been published about behaviors of Internet traffic in Thailand. We cannot rely on overseas Internet reports to make planning decisions on local networks. For one reason, the Thai telecom industry has a different structure than that of other countries. Moreover, different supporting infrastructure gives different contexts in which users' behaviors are shaped upon. These reasons motivate the need to study behaviors of Internet traffic in Thailand.

The Thai telecom industry is governed by strict regulations. Such regulations play a major role in dictating the growth in data networks. As a result, many reports about Internet traffic patterns, Internet growths, and their implications in one country may not be applicable to Thailand. For example, the U.S. Internet traffic growth has been through a period of dot-com boom, burst, and now recovery. Although its traffic approximately doubles each year, U.S. operators sees no need for new fiber deployment in the near future [9]. Yet such cycle has not been experienced by the Thai telecom industry. Therefore, we cannot possibly argue the same implications or make the same forecast even if our Internet traffic has been growing at a similar rate.

This paper reports on traffic characteristics of the Internet usage in Thailand by looking at aggregate link statistics collected via the Simple Network Management Protocol (SNMP). We report on both short-term and long-term traffic variability, as well as comparing the characteristics of domestic and international traffic. From these characteristics, it is clear that the traffic has both daily and weekly periodic components as well as a longer-term growth trend.

This paper is organized as follows. Section 2 outlines previous studies in the area of Internet traffic measurements. Section 3 describes the traffic data used in our analyses. Section 4 explains various patterns and characteristics of Internet traffic in Thailand. Section 5 discusses implications of the findings. Finally, the paper concludes in Section 6.

2. PREVIOUS STUDIES

Internet traffic measurement has been a subject of interest in the past decade. However, there have been a limited number of publications on large-scale studies of Internet traffic. This is probably due to the competitive nature of commercial Internet providers. Traffic and utilization data is usually guarded as highly confidential.

Traffic measurements on Internet backbone are mostly reported by tier-1 carriers, such as MCI [1], Sprint [2][3], and AT&T [4][5]. These measurements generally use statistics collected from SNMP management objects, as well as packet-level and flow-level traces.

Previous studies of Internet traffic can be categorized into several groups: those that look at a long-term trend for forecasting purposes [3][9], those that look at traffic variability in a small time scale [4][5], those that focus on just the "busy hour" [6], and those that investigate traffic compositions [1]. In this paper, we will focus on both short-term and long-term variability because we want to get a big picture of Internet behaviors in Thailand.

For Thai Internet, currently the Internet Information Research (IIR) website (http://iir.ngi.nectec.or.th) maintained by NECTEC is the main source of statistics on Internet growths and traffic measurements. In corporation with CAT Telecom Plc., the IIR website obtains SNMP measurements at domestic and international Internet exchange gateways, and reports non-business sensitive statistics. Currently the measurements are limited to SNMP data only. (Packet-level measurements are not available.) However, no extensive analysis has been made using this traffic data.

To our knowledge, this article is the first to report and analyze characteristics of Internet traffic in Thailand based on actual SNMP measurements at both domestic and international exchange gateways.

3. TRAFFIC DATA

The traffic data used in this study comes from three major Internet exchanges in Thailand—Internet Information Research (IIR), National Internet Exchange (NIX), and International Internet Gateway (IIG). Presently, all Internet service providers (ISPs) in Thailand are required to

exchange their domestic traffic at NIX, and their international traffic at IIG. IIR mainly serves as local exchange for academic, research, and government networks. Some ISPs may connect to IIR as well. (See Fig. 1.) We collect SNMP incoming and outgoing byte counts for all interfaces at these three exchange locations. The byte counts were polled every 5 minutes from July 2003 until November 2005.

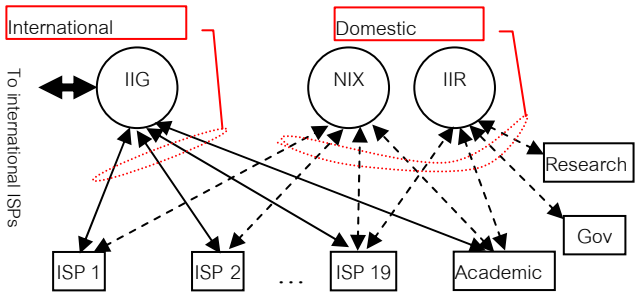

Fig. 1: *Logical Internet connectivity in Thailand. For a detailed map, see [http://www.nectec.or.th/internet/map/]*

To obtain the incoming and outgoing bit rate at each interface, we rely on the Multi Router Traffic Grapher (MRTG) tool [7]. For short-term statistics, MRTG generates a bit rate from the byte volume over a 5-minute interval. In addition, MRTG manages a year long traffic trace by collecting daily 5-minute averages, weekly 30-minute averages, monthly 2-hour averages, and yearly 1-day averages.

Our traffic traces have a few limitations. First of all, data may be missing because SNMP messages use UDP transport, and there are brief periods where our storage server went down. Secondly, the data may be incorrect due to poor router implementation or MRTG configuration. (MRTG could fill in the missing SNMP data with previous measurement data or use 0 to represent the missing data.) In this paper, we ignore these flaws in the measurements since we are only interested in the overall aggregated trends and behaviors, not in the precise statistics of the small-scale variability.

We obtain 118 traces of incoming and outgoing bit rates from 59 distinct interfaces connected to IIR, NIX and IIG. Note that some ISPs may have more than one interface at each exchange point. In that case we aggregate the bit rates of these interfaces into one trace per ISP.

4. TRAFFIC CHARACTERISTICS

For the rest of this paper, the term "domestic traffic" refers to traffic going to and leaving from both NIX and IIR gateways, without making distinction between the two

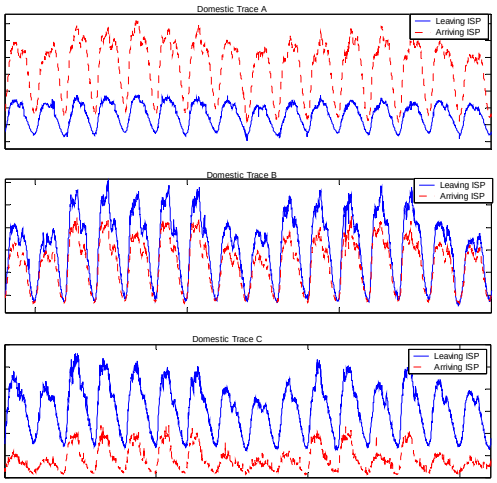

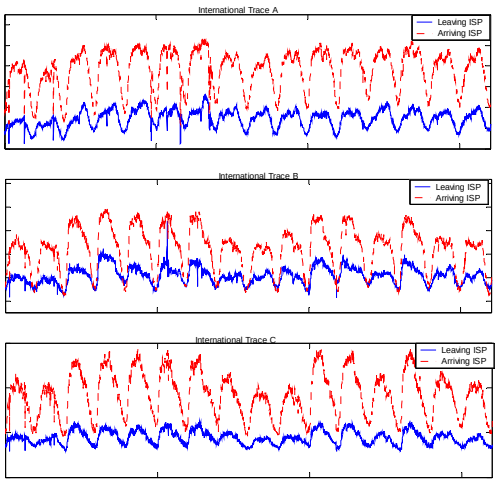

Fig. 2 *(a) Domestic traces during May 7 – 23, 2005*

(b) International traces during April 23 – May 9, 2005

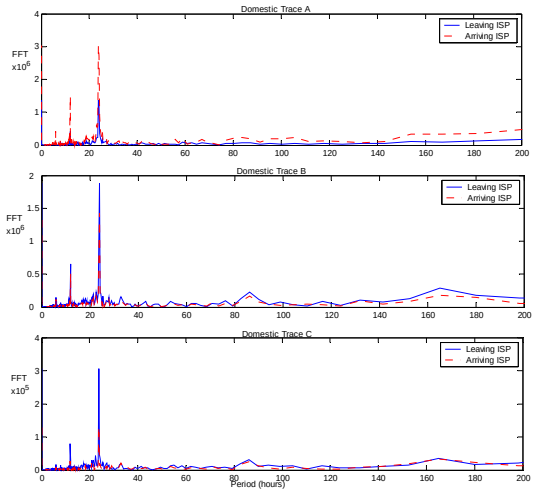

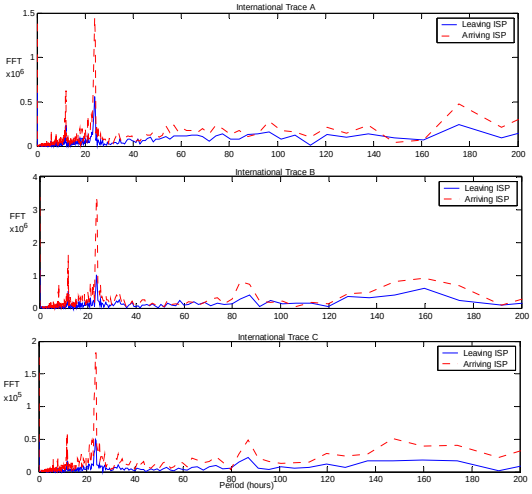

Fig. 3 *(a) Fast Fourier transform for domestic traces*

(b) Fast Fourier transform for international traces

locations. Likewise, the term "international traffic" refers to traffic going into and out of IIG on the ISP side. Therefore, when we talk about "international capacity" or "domestic capacity," we refer to capacity of communication lines connecting ISPs to IIG, or ISPs to NIX and IIR, respectively.

4.1 Short-Term Behaviors

Let us first examine the traffic patterns during a two-week period. Six sample traces are shown in Fig. 2. Similar to previous reports on Internet traffic patterns [1][3][4][5], we found strong diurnal (daily) and weekly cycles across all traces at all three exchange locations. The daily cycles rise and fall synchronously for both incoming and outgoing directions.

It is clear that there are strong periodicities in the data. In order to verify their existence, we calculate the Fourier transform for all traces, with six examples shown in Fig. 3. The results indicate the dominant period at 24 hours across all traces. The 12-hour period is also present to a lesser degree. The weekly cycle (168-hour period) manifests itself only in some traces. When the weekly period exists, it is usually weaker than the 12-hour period.

The fact that the 12-hour period is quite dominant in our traffic data is interesting, albeit expected. Previous studies by the U.S. tier-1 carriers never show such evident 12-hour dip in their traffic [1][3][5]. This is probably because the traffic measurements in the U.S. span across three different time zones, while the Thai measurement spans across just one time zone.

4.2 Domestic vs. International Traffic

In this section we compare patterns of domestic traffic and international traffic. The daily and weekly patterns are similarly evident for both types of traffic. However, one noticeable difference lies in the volume of incoming and outgoing traffic. Here, incoming traffic refers to traffic that goes toward ISPs, and outgoing traffic refers to traffic that goes from ISPs toward the exchange gateways.

For domestic traffic, the volume of traffic in both directions differs for different ISPs. However, for international traffic, all traces exhibit higher (or equal) incoming traffic volume than outgoing traffic volume. This finding is not surprising since it is more likely for Thai users to download large files from servers abroad than to upload them.

4.3 Long-Term Trends

Next we examine the long-term trend of Internet traffic in Thailand. Fig. 4 plots the growth of domestic and international capacity in comparison to the growth of actual traffic demand. The capacity data is from July 1997 to February 2006. However, the measurements of actual traffic demand, which is the aggregation of demands from all interfaces, are only available from December 2003 to November 2005.

At domestic exchanges, the capacity deployment remains stable with very little activity up until end of 1999. After that the capacity seems to explode. For example, during the year 2003, the rate of growth is as high as 4x. However, capacity growth appears to have declined to about 2x during 2004 and 1.5x during 2005.

The domestic Internet traffic, however, does not increase as fast as the total capacity. The traffic demand increases steadily at about 2x per year. This indicates a high degree of over-provisioning in the domestic Internet. In addition, we point out that the total incoming traffic demand has the same magnitude as that of outgoing traffic demand. This is expected since traffic cannot disappear inside the domestic exchanges.

Next let us look at the trend of international capacity deployed by ISPs. Unlike that of domestic capacity, international capacity has been growing at a steadier rate—approximately doubling annually since 1997. In addition, we see asymmetric capacity deployment during 1999-2002. This is a result of many ISPs expanding their lines using "simplex" satellite service in order to cope with congestions in the inbound traffic. In late 2002, most ISPs opt for high-quality fiber links and symmetric satellite services. This results in equality between the inbound and outbound capacities [8].

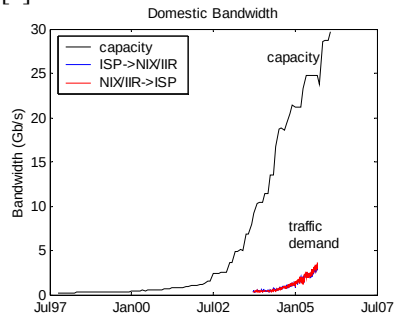

(a) Domestic bandwidth

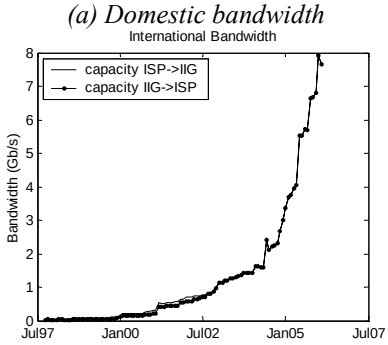

(b) International bandwidth

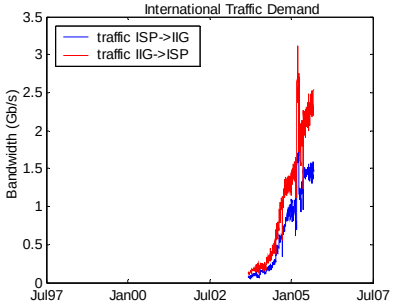

(c) International traffic demand

Fig. 4: *Internet capacity in comparison to traffic demand*

The international traffic demand, however, appears to grow faster than 2x annually. The average rate of increase over the two-year period is almost 4x per year (although a close examination reveals a faster growth rate during the first

year than during the second year). Moreover, the rate of growth is similar for both incoming and outgoing traffic, but the traffic coming from abroad consumes more bandwidth than the traffic leaving Thailand. This confirms previous observations we made on a small time scale.

Note that, although the growth rate of international traffic is higher than that of domestic traffic, the amount of traffic volume at both locations is quite comparable, on the order of 3 Gb/s daily.

5. IMPLICATIONS

This section summarizes the main findings and discusses their implications.

(1) As expected, Internet traffic exhibits strong daily and weekly periodicities. This suggests predictability of the traffic data. One can always expect a higher transfer rate when using Internet on weekends or during midnight-6am on weekdays. As a result, large file transfers or bulk transactions should be planned during these periods.

(2) Volume of Internet traffic coming from abroad is larger than that leaving Thailand. This suggests traffic deficit with foreign ISPs. As a result, Thai ISPs will be at disadvantage when dealing with overseas peering agreements.

(3) Domestic bandwidth usage is low. Only 10% of capacity deployed by ISPs is used daily. Although such usage level is considered normal for large Internet back-bones in the U.S. [9], it may not be appropriate for Thai networks. On a pessimistic side, such low utilization level could mean that there is a low demand for local Internet contents, or there are not enough interesting local contents available. Both reasons could worsen the degree of overseas traffic deficit mentioned above. On an optimistic side, light utilization is partially a reflection of the high growth rate. A low utilization level may signal an era of competition and overcapacity. It is speculated that such over-provisioning situation is driven by the desire for low transaction latency more than the desire for actual data transfer [9].

(4) Deployment of international capacity is slower than that of domestic capacity while international traffic demand grows faster than domestic traffic demand. This is due to following reasons. Currently, there are many limitations to ISP's deployment of international capacity. Firstly, regulations require that ISPs obtain international bandwidth through a state-owned enterprise. In other words, they cannot peer with their overseas counterpart directly. Secondly, the declining value of the Thai Baht after the 1997 crisis puts Thailand at a disadvantage position for international bandwidth trade. Both reasons amount to high cost of provisioning new international capacity.

As for why traffic demand grows faster for international traffic than domestic traffic, this could be attributed to the user's insatiable demand for international contents. This, in turns, implies that international bandwidth is essentially driven by the desire for actual network usage, as oppose to the desire for low transaction latency as seen in domestic traffic.

(5) Fast Internet growth rate found in Thailand implies that careful capacity planning is no longer possible. Even a rate of traffic doubling each year is considered disruptive for network operators, let alone the 4x growth [9]. The growth of Internet in Thailand is not likely to be smooth or predictable. Thus, simplicity should be the key to network engineering. Over-provisioning might be the way to go since it can reduce the overhead cost for putting in more bandwidth every so often.

6. CONCLUSIONS

This paper studies behaviors of domestic and international Internet traffic in Thailand. The observations from the SNMP measurements lead to important findings and insights into the current state of Internet in Thailand.

Although Internet in Thailand appears to grow at a normal, healthy rate, we believe that it is still far from reaching its full capacity. Many limitations exist due to the current telecom regulations. We expect to see a dramatic change in the growth of Internet in Thailand in a near future when the telecom industry is liberalized. When more exchange points are set up, and ISPs are permitted to interconnect with overseas peers more freely, competitive pricing would take place, which in turn stimulates further explosion in the growth rate.

Our future work is to exploit these findings in the modeling and forecasting processes. Moreover, we plan to use the same set of measurements to analyze characteristics of overload and anomalous traffic events.

7. ACKNOWLEDGEMENT

I would like to thank CAT Telecom Plc. for allowing access to SNMP data at the NIX and IIG gateway routers, and Phusit Roongroj and Chatchai Chan-In for their work in collecting and managing the traffic data used here.

REFERENCES

[1] K. Thompson, G. Miller, R. Wilder, "Wide-Area Internet Traffic patterns and Characteristics," *IEEE Network*, Nov/Dec 1997.

[2] Sprint Advanced Technology Lab, Publications website, [http://ipmon.sprintlabs.com/pubs_trs/webpages/ATLPUB.php].

[3] K. Papagiannaki, N. Taft, Z. Zhang, C. Diot, "Long-Term Forecasting of Internet Backbone Traffic: Observations and Initial Models," *IEEE Infocom*, Apr. 2003.

[4] P. Pongpaibool, R. Doverspike, M. Roughan, J. Gottlieb, "Handling IP traffic surges via optical layer reconfiguration", *OFC2001*, Mar. 2001.

[5] M. Roughan, A. Greenburg, C. Kalmanek, M. Rumsewicz, J. Yates, Y. Zhang, "Experience in Measuring Internet Backbone Traffic Variability: Models, Metrics, Measurements and Meaning," *ITC '2003*, Sept. 2003.

[6] T. Mori, R. Kawahara, S. Naito, "On the Characteristics of Internet Traffic Variability: Spikes and Elephants," *IEICE Trans. Inf. & Syst.* Vol. E87-D, No. 12, Dec. 2004.

[7] T. Oetiker, D. Rand, "MRTG: The Multi Router Traffic Grapher," [http://people.ee.ethz.ch/~oetiker/webtools/mrtg/].

[8] NECTEC's *Thailand ICT Indicators 2005*, February 2005, pp. 34-35.

[9] A. Odylzko, "Internet traffic growth: Sources and implications," *Proc. SPIE*, vol. 5247, 2003, pp. 1-15.

[10] S. Palasri, S. Huter, Z. Wenzel, *The History of the Internet in Thailand*, University of Oregon Books, 1999.