Dynamics of the IP Time To Live Field in Internet Traffic Flows

Sebastian Zander, Grenville Armitage, Philip Branch
Centre for Advanced Internet Architectures, Technical Report 070529A
Swinburne University of Technology
Melbourne, Australia
szander@swin.edu.au, garmitage@swin.edu.au, pbranch@swin.edu.au

***Abstract*—The Time To Live (TTL) field present in the IP protocol header is used to limit the lifetime of packets in the network. Previous research has measured TTL for studying path lengths and dynamics in IP networks, and for detecting route changes. How the TTL varies over short timescales of subsequent packets of traffic flows has not yet been analysed. Such knowledge is needed for passively detecting route changes based on existing traffic in the network or traffic traces, and for designing mechanisms that modulate the TTL field such as IP traceback techniques and covert channels. In this paper we analyse small time scale TTL variation in Internet traffic flows based on a number of packet traces captured at different locations in the Internet.**

***Index Terms*—Network Measurement, IP Time To Live, Small Time Scale Dynamics**

I. Introduction

The IP Time To Live (TTL) header field limits the lifetime of an IP packet, preventing packets from living forever during routing loops [1]. A packet's initial TTL is set by the sender and then decremented by each network element along the path processing the packet's IP header (e.g. routers and firewalls). Packets are discarded if their TTL becomes zero while still in transit.

Initial TTL values used by popular operating systems are well known [2], [3], and the number of hops (usually referred to as hop count) between sender and receiver (or measurement point) can be estimated by subtracting the observed TTL from the closest initial TTL value. Previous research has done this to study the size of the Internet in hops and its dynamics (IP path changes) over time periods of several days up to several months [4], [5], [6], to investigate the predictive power of TTL for finding lowest latency paths [7] or to detect route changes [8].

However, previous research has not analysed how TTL varies in small time scales of subsequent packets of Internet traffic flows (sequences of packets with the same IP addresses, ports, and protocol). The motivation for investigating this is twofold. Firstly, for passively studying IP path changes based on existing traffic in the network or captured traffic traces (instead of active sampling paths as it was done previously) one needs to understand the relation between TTL changes and path changes. If TTL changes are not only caused by path changes one cannot simply infer path changes from TTL changes and needs to differentiate TTL changes. Secondly, this research provides valuable information for people developing techniques that modulate the TTL field, such as IP traceback techniques [9] or covert channels [10], [11].

We study all TTL changes occurring in consecutive packets of traffic flows, instead of purely focusing on TTL variation caused by IP path changes. Based on a number of packet traces captured at different locations in the Internet we analyse the characteristics of TTL changes and describe the most common change patterns observed. Our main findings are:

- The amount of TTL variation is fairly small: in general $< 1\%$ of the packet pairs and $< 6\%$ of the flows experience TTL changes.
- Most flows with TTL changes have only 2 distinct TTL values and the hop count difference is generally ≤ 3.
- Most flows have only few transitions between different TTLs, but there are flows with periodic TTL changes and frequent seemingly random TTL changes.
- We cannot find a strong correlation between the number of TTL changes and the estimated hop count, which is probably because most TTL changes are unrelated to routing and apparently caused by middleboxes such as firewalls and proxies.

The paper is organised as follows. In section II we review related work. In section III we describe the datasets and the methodology used for extracting the TTLs. Section IV presents the results of the analysis and Section V describes the most common TTL change patterns. Section VI concludes and outlines future work.

II. Related Work

Fei et al. measured hop counts and round trip times (RTTs) from one host to a random set of target hosts [4]. They found that in the US most hosts are within a range of 18 hops and although RTT increases with the number of hops there is no strong correlation with hop count.

Begtasevic and Mieghem measured the hop counts between one host and several thousand destinations distributed across the world over a time period of one week roughly every hour [5]. During their measurement period the hop counts towards a significant number of destinations changed. Even when the hop count value remained constant often the path from source to destination had changed.

Huffaker et al. studied the change of the hop count distribution from a source to a set of destinations over several months and compared the distributions obtained from multiple sources [6]. The hop count distribution was similar for most sources and did not significantly change over time.

Huffaker et al. also studied the predictive power of hop count (IP path length) in the selection of the lowest latency destination from a set of alternatives [7]. They found that hop count is a better predictor than Autonomous System (AS) path length, but worse than geographic distance or RTT.

Shen et al. investigated if the TTL field can be used to detect route changes [8]. According to their study many route changes can be detected, but TTL monitoring alone is inadequate for a very reliable detection.

In contrast to the previous work, we study the characteristics of all TTL changes – not only those related to path changes – occurring within the small time scale of consecutive packets of Internet traffic flows. The purpose of our study is to identify the main factors that cause TTL changes in packet flows and analysing the change characteristics.

III. Datasets and Methodology

In this section we describe the different datasets we use and how we extracted TTL data from them.

A. Datasets

We use packet traces of different size, origin and date for our analysis (see Table I). The CAIA trace contains only game traffic, the Grangenet trace contains game and web traffic, and the other traces contain a mix of traffic (including web, peer-to-peer, game, and email traffic). For the CAIA and Grangenet traces we have removed all flows originating from the game servers, because their TTL is constant (initial TTL).

B. Packet Flows and Packet Pairs

Usually analysis of network traffic is based on packet flows (sequences of packets sharing common source/destination addresses, port numbers, protocol over some time period). Because we found the number of TTL changes is correlated with the number of packets and the duration of flows, we decided to analyse the characteristics of TTL changes based on packet pairs (defined as two subsequent packets of a flow as shown in Figure 1). This isolates the characteristic under study (e.g. estimated hop count) from correlated flow properties (e.g. size and duration). However, for analysis requiring a longer sequence of packets we use packet flows (e.g. Section IV-E).

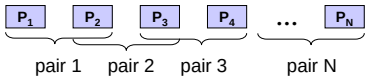

Figure 1. Relation between packet flows and packet pairs

C. TTL Data Extraction

Packets are grouped into flows according to IP addresses, port numbers and protocol. The end of a flow is determined by TCP connection teardown or a 600 second idle timeout (whichever occurs first). Although we treat flows as bi-directional for the TCP state tracking and flow timeout, we later separate the two directions into unidirectional flows because TTL variation is different in either direction. In the following the notion of flow is always unidirectional. We only consider flows with at least four packets and an average packet rate of at least one packet per second.

For each flow we extract the series of TTL values from the IP packet headers as observed at the capture location. This is also a list of TTL pairs as described in the previous section. If different TTL values occur in a flow or packet pair this is referred to as TTL change. Otherwise we refer to the TTL being constant.

Table I
PACKET TRACES USED IN THE ANALYSIS

Trace	Capture Location	Date	Public
CAIA	Public game server at Swinburne University in Melbourne, Australia	05/2005 - 06/2006	no
Grangenet	Public game server connected to the Grangenet research network in Canberra, Australia [12]	05/2005 - 06/2006	no
Twente	Aggregated uplink of a Asymmetric Digital Subscriber Line (ADSL) access network [13]	07/02/2004 - 12/02/2004	yes
Leipzig	Internet access router at a large German university [14]	21/02/2003	yes
NZIX	New Zealand exchange point [14]	05/07/2000 - 06/07/2000	yes
Bell	Firewall at Bell Labs [14]	19/05/2002 - 20/05/2002	yes
Waikato	University of Waikato [15]	04/05/2005	on request

D. Effective Path Sampling Rate

We have chosen the average packet rate instead of a maximum inter-arrival time threshold as flow selection criterion because it is less restrictive and provides us with more data. Using the average raises the question of how often paths were sampled. Figure 2 shows a plot of the distributions of the 0.95-quantile of the packet inter-arrival times.

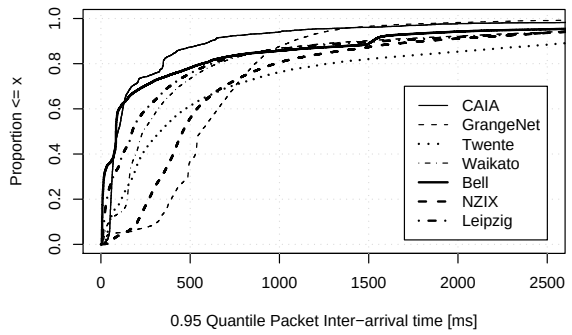

Figure 2. Distribution of the 0.95-quantile of packet inter-arrival times of flows with TTL changes

Across all traces approximately 80% of the packet pairs have inter-arrival times of ≤ 1 second meaning 80% of the packet pairs sampled the paths at least once every second.

IV. ANALYSIS OF TTL VARIATION

A. Amount of TTL Variation

Table II shows the number and the percentage of flows and packet pairs with and without TTL changes (absolute numbers are rounded to kilo or mega flows).

Overall, we find the TTL is constant for the majority of flows, but a small number of flows do experience TTL changes. The percentage of flows with TTL changes is between 2-6%, except for the NZIX trace which has a significantly larger number of flows with TTL changes. The percentage of packet pairs with TTL change is between 0.02-0.5%, except for NZIX where it is significantly higher (5.1%). The very small percentages show that if there is TTL variation in flows, changes occur only for a very small number of the flows' packet pairs.

B. Distinct TTL Values per Flow

Figure 3 shows the distribution of the number of distinct TTL values per flow (note the y-axis scaling).

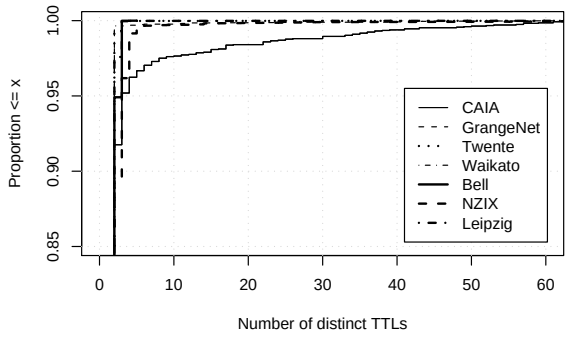

Figure 3. Distribution of the number of distinct TTL values per flow

Flows with TTL variation mostly have only two distinct TTL values. Only 10% and less of the flows have more than two values and flows with more than five different TTLs are very rare, except in the CAIA trace. In the CAIA trace a number of flows encountered what seems to be routing anomalies (see section V).

C. TTL and Hop Count Change Amplitudes

Figure 4 shows the amplitude of the TTL change of packet pairs. We define the amplitude as the difference between the maximum and the minimum TTL value (peak-to-peak amplitude).

Table II
FLOWS AND PACKET PAIRS WITH/WITHOUT TTL CHANGES AND PERCENTAGE OF CHANGES.

	Flows w/o change [k]	Flows with change [k]	Flows with change [%]	Packet pairs w/o change [M]	Packet pairs with change [k]	Packet pairs with change [%]
CAIA	128.6	2.8	2.1	1,456.3	340.0	0.02
Grangenet	283.0	8.6	2.9	215.7	62.1	0.03
Twente	1,354.6	24.7	1.8	95.5	74.8	0.08
Waikato	1,255.9	57.8	4.4	21.2	86.4	0.4
Bell	899.8	52.9	5.6	36.4	87.1	0.2
NZIX	2,482.7	510.8	17.1	91.1	5,039.5	5.5
Leipzig	7,155.1	429.1	5.7	365.5	1,822.7	0.5

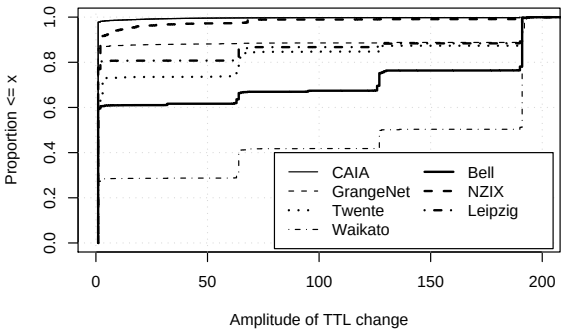

Figure 4. Distribution of the amplitude of TTL change of packet pairs

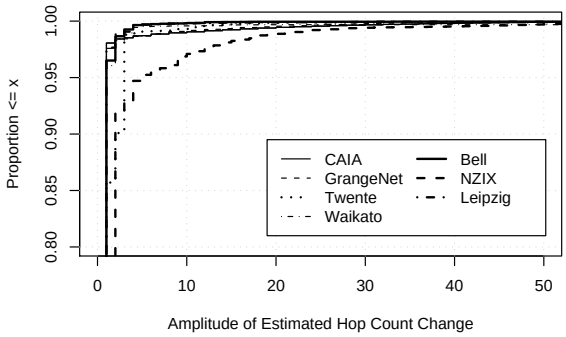

Figure 5. Distribution of the estimated hop count amplitude of packet pairs

The amplitude is one for many packet pairs, but in some traces large numbers of packet pairs have amplitudes around 64, 127, or 191. The reason for these high amplitudes is middleboxes (firewalls or proxies) sending packets as part of the TCP handshake or teardown on behalf of clients (e.g. for SYN flood attack protection). The TTLs in these packets are set to the initial TTL of the firewall/proxy, which can differ from the initial TTL used by the host behind the firewall. Different operating systems use different initial TTLs, the most common being: 64 (Linux, FreeBSD), 128 (Windows 98/XP), 255 (Cisco) [2], [3]. Therefore, the difference of two TTLs will be 64, 127, or 191 plus or minus the distance between middlebox and host. The CAIA trace does not contain any TCP traffic and therefore does not have high amplitudes.

Figure 5 shows the amplitude of the estimated hop counts of packet pairs. This amplitude is defined as the difference between estimated maximum and minimum hop counts. The hop count is estimated by subtracting a packet's TTL value from the closest initial TTL. We assume the initial TTL is 32, 64, 128 or 255 (these values cover all currently popular operating systems, see [2], [3]).

For most packet pairs with TTL changes the hop count changes only by one (two for NZIX and three for Twente). Only a small number of packet pairs have larger hop count changes (NZIX has the highest percentage).

D. Correlation of TTL Changes and Hop Count

Figure 6 shows the estimated hop count distribution of packet pairs without TTL changes (hop count was estimated as described in the previous section).

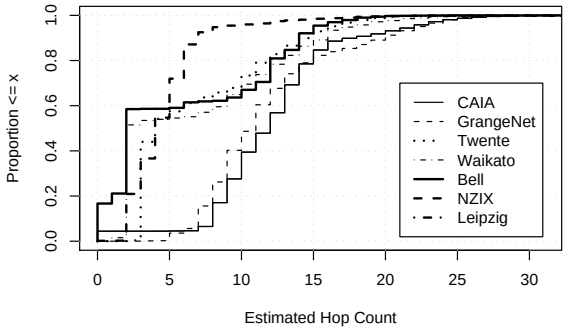

Figure 6. Distribution of estimated hop count of packet pairs without TTL change

The CAIA and Grangenet traces show a similar distribution with no low hop count values because server

to client traffic was removed (see Section III-A). The Leipzig, Twente, Waikato and Bell traces were captured at access links zero to three hops away from the local hosts. The upper half of the hop count distribution is very similar for all these traces, except for NZIX that shows a tendency towards smaller hop counts (probably caused by more localised traffic because NZIX is a peering point between New Zealand Internet service providers).

We also estimated the hop count of packet pairs with TTL change by subtracting the larger of the two TTLs from the initial TTL (see Figure 7).

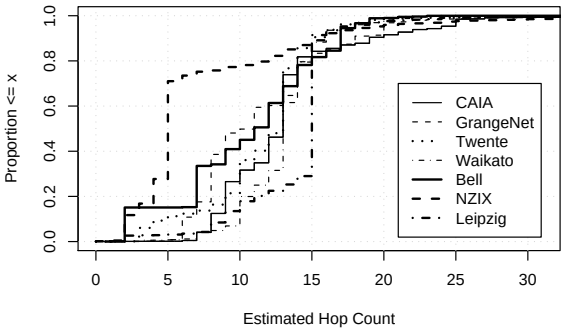

Figure 7. Distribution of estimated hop count of packet pairs with TTL change

Comparing Figure 6 and Figure 7 it is clear that almost none of the local packet pairs originating from three or less hops away experienced a TTL change. However, the distributions of packet pairs with larger estimated hop counts look similar in both figures.

We also computed Pearson's correlation coefficient between the estimated number of hops and the number of TTL changes per flow. The result indicates that the amount of TTL changes per flow is not linearly correlated with the estimated number of hops.

E. Frequency of TTL Changes

Figure 8 shows the number of TTL changes per flow (x-axis limited at 20 changes).

For most datasets the majority of flows only has few TTL changes, but a large percentage of flows in the CAIA trace has a large number of changes. This is caused by the fact that the CAIA trace only contains long UDP flows of which many have periodic TTL changes (see Section V), but TCP flows with a fixed number of changes are predominant in all other traces.

Figure 9 depicts the change frequency for flows with at least six TTL changes. The TTL change frequency of a flow is defined as the number of TTL changes divided by the total number of packet pairs.

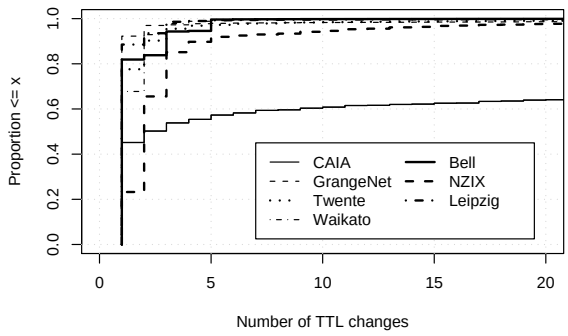

Figure 8. Distribution of number of TTL changes per packet flow

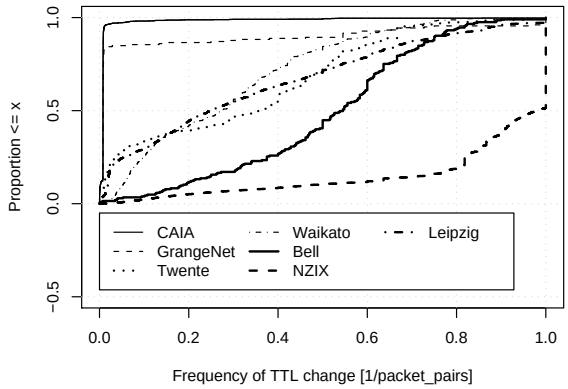

Figure 9. Distribution of frequency of TTL changes for flows with at least six TTL changes

CAIA and Grangenet traces have very low change frequencies because most flows have infrequent periodic TTL changes. Twente, Leipzig, Waikato and Bell have higher change frequencies, with roughly half of the flows changing TTL on average every third to second packet pair. While the distributions of Twente, Leipzig and Waikato are fairly similar, Bell has higher change frequencies. In the NZIX trace in most flows TTL changes at least every second packet pair (presumingly caused by load balancing or route flapping).

F. Probability of Uncommon TTLs in Flows

Figure 10 shows the average empirical probability that the *n*-th packet of a flow has a TTL value differing from the most common TTL value occurring in the flow (limited to the first 50 packets). This probability is computed as number of packets not having the most common TTL value divided by the total number of packets for each packet position *n* averaged across all flows. Since most flows are short the number of samples decreases with increasing *n*.

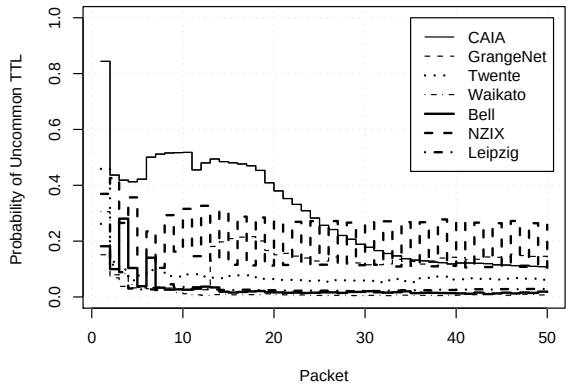

Figure 10. Probability of uncommon TTL values depending on the position of the packet in the flow (up to the first 50 packets)

The probability for an uncommon TTL is clearly not uniformly distributed. It is usually the highest at the flow start and then decreases. For some datasets there is a little peak at the end of typical short TCP flows (4-6 packets). CAIA and Grangenet datasets have a high probability until 20-30 packets caused by game flows where the first up to 30 packets at the start of the flow have a different value. For NZIX the probability oscillates between two values, which is caused by the flows with very frequent TTL changes in this trace.

V. Change Patterns

In this section we describe the most common change patterns and their most likely causes. We group the patterns into four main classes:

- Deterministic: TTL changes occurring in specific packet pairs of a flow (e.g. at the start of flows).
- Periodic: TTL changes occurring periodically every *n* packets.
- Infrequent random: random TTL changes occurring rarely (e.g. route changes or anomalies).
- Frequent random: random TTL changes occurring frequently (e.g. load balancing, route flaps).

Obviously, a single traffic flow can experience multiple different change patterns.

A. Deterministic Changes

In many UDP and TCP flows the first packet has a TTL value one less than the common value (hop count one higher). We suspect multi-layer switching devices that do not adjust the TTL could cause this. In multi-layer switching only the first IP packet in a flow is routed while the rest are switched.

Many TCP flows have TTL changes, where the TTLs differ by more than 32. This is caused by middleboxes (e.g. firewalls, proxies) that send some packets part of the TCP handshake and teardown on behalf of the client. In case the middlebox's initial TTL differs from the host's initial TTL the characteristic amplitudes are 32, 64, 191 plus or minus the number of hops between the host and the middlebox. Otherwise the difference is just the number of hops the middlebox is away from the host (1-3 hops). More evidence that these packets originate from different devices is that they sometimes also differ in other IP header fields e.g. the Don't Fragment bit.

In some UDP and TCP flows the TTL decreases by one every packet. We suspect this to be some kind of (TCP) traceroute.

In some UDP flows the first (few) packet(s) have a TTL one higher than the common value (hop count one lower). Often this effect occurs in combination with periodic TTL changes (see section V-B).

B. Periodic Changes

We found a number of UDP and some TCP flows have periodic TTL changes (confirmed by the auto-correlation). In some flows every 256-th packet has a TTL one higher than the majority of TTLs. In many cases this is exactly every 256-th packet over long time periods, but sometimes the period is not exact and varies by few packets. Furthermore, sometimes instead of exactly one packet a short burst of packets with TTLs one higher occurs.

We are not sure what causes this pattern, but it seems not specific to some country or service provider. This could mean devices close to the end users such as firewalls or access routers cause it. The vast majority of long UDP game flows with TTL changes in the CAIA and Grangenet traces have this pattern, and it also occurs in some long UDP and TCP flows of the Twente trace.

C. Infrequent Random Changes

In a number of flows the TTL value is constant for a large number of packets, then it changes to a different value and stays constant for a large number of packets again. Often only one of these changes occurs in a flow. We suspect these changes are caused by route changes.

Few flows experience drastic changes of TTL values within 10-20 consecutive packets (observed mostly in the CAIA trace). The TTL changes by large amounts (up to 30) between consecutive packets and usually spirals downwards to very low TTL values, before returning to the normal value. All the packets comprising these anomalies arrive compressed (back to back within microseconds) although UDP game packets are spaced

milliseconds apart at the sender. We are not sure what causes this behaviour, but it appears the packets are involved in some kind of short-term routing anomaly.

D. Frequent Random Changes

Some UDP and TCP flows have very fast, random TTL changes (TTL values change almost every packet), which is presumingly caused by IP load balancing or fast route flapping. Some TCP flows show a random mix of two TTL values that differ by more than 32 indicating that packets are sent from devices with different initial TTL. We believe application-layer load balancing or proxying coud cause this.

VI. Conclusions

In this paper we analysed changes in the TTL values between subsequent packets of traffic flows based on several traffic traces captured at different locations in the Internet. We found that TTL changes are not very common but appear in a significant part of the data across all traffic traces. Most of the changes are apparently not caused by routing changes, but by middleboxes such as firewalls and proxies. We have analysed the various characteristics of TTL changes and described the main change patterns observed.

There are a number of issues left for future research. The analysis should be extended towards further traffic traces. Also there are some metrics we have not yet analysed, for example the randomness and burstiness of TTL changes. If unanonymised datasets were available, the correlation between AS path length or geographical location and TTL changes could be analysed. Furthermore, we plan to develop methods for quantifying how often the different common change patterns occur in the data. We think such classification could benefit from combining TTL change information over multiple parallel (time and spatial) packet flows. Finally, an in-depth analysis of middlebox induced change patterns could be used for developing fingerprinting techniques for such devices.

References

[1] J. Postel, "Internet Protocol," RFC 0791, IETF, Sept. 1981. http://www.ietf.org/rfc/rfc0791.txt.

[2] C. Jin, H. Wang, K. Shin, "Hop-Count Filtering: An Effective Defense Against Spoofed DoS Traffic," in *Proceedings of the 10th ACM International Conference on Computer and Communications Security (CCS)*, pp. 30–41, October 2003.

[3] N. Davids, "Initial TTL Values." http://members.cox.net/~ndav1/self_published/TTL_values.html.

[4] A. Fei, G. Pei, R. Liu, L. Zhang, "Measurements on Delay and Hop-Count of the Internet," in *Proceedings of IEEE GLOBECOM - Internet Mini-Conference*, 1998.

[5] F. Begtasevic, P. van Mieghen, "Measurements of the Hop Count in the Internet," in *Proceedings of Workshop on Passive and Active Measurement (PAM)*, pp. 183–190, April 2001.

[6] B. Huffaker, M. Fomenkov, D. Moore, K. Claffy, "Macroscopic analyses of the infrastructure: Measurement and visualization of Internet connectivity and performance," in *Proceedings of Passive and Active Measurement Workshop (PAM)*, April 2001.

[7] B. Huffaker, M. Fomenkov, D. J. Plummer, D. Moore, k claffy, "Distance Metrics in the Internet," in *Proceedings of the IEEE International Telecommunications Symposium (ITS)*, September 2002.

[8] C. Shen, H. Schulzrinne, S. Lee and J. Bang, "Routing Dynamics Measurement and Detection for Next Step Internet Signaling Protocol," in *Proceedings of IEEE/IFIP Workshop on End-to-End Monitoring Techniques and Services (E2EMON)*, May 2005.

[9] E. Jones, O. Le Moigne, J.-M. Robert, "IP Traceback Solutions Based on Time to Live Covert Channel," in *Proceedings of 12th IEEE International Conference on Networks (ICON)*, pp. 451–457, November 2004.

[10] S. Zander, G. Armitage, P. Branch, "Covert Channels in the IP Time To Live Field," in *Proceedings of Australian Telecommunication Networks and Applications Conference (ATNAC)*, December 2006.

[11] S. Zander, G. Armitage, P. Branch, "A Survey of Covert Channels and Countermeasures in Computer Network Protocols," *(accepted for publication in) IEEE Communications Surveys and Tutorials*, 2007.

[12] Grid and Next Generation Network (GrangeNet). http://www.grangenet.net/.

[13] R. van de Meent, "M2C Measurement Data Repository," December 2003. http://m2c-a.cs.utwente.nl/repository/.

[14] NLANR PMA: Special Traces Archive. http://pma.nlanr.net/Special/.

[15] R. Nelson, "WAND Research Group." http://www.wand.net.nz/index.php.