

On the Session Lifetime Distribution of Gnutella

Kevin Ho¹, Jie Wu² & John Sum³

¹Department of Computer Science and Communication Engineering
Providence University, Sha-Lu, Taichung, Taiwan ROC.

²Department of Computer Science and Engineering, Florida Atlantic University
Boca Raton, Florida 33431, USA.

³Department of Information Management, Chung Shan Medical University
Taichung 402, Taiwan
(01/2007)

Over the past few years, capturing the characteristics of the Gnutella overlay network has been one of the major research activities in distributed computing. While most of the works are focusing on the node degree distribution, upstream/downstream traffics and file distribution, only a few have been done on modeling the session lifetime distribution. A good model of the session lifetime distribution is basically a key to understand the formation and evolution of the topological structure of Gnutella. In this paper, the main objective is to present a measurement result we did in 2003 on the session lifetime distribution and make comparison with the results obtained by other researchers during 2001 to 2005. In accordance with our measurement, it is found that the lifetime distribution fits to a power law distribution with exponential cut-off. Specifically, for all $t \geq 15$ minutes, $P_l(t) \propto (t/780)^{-1.12} \exp(-t/780)$. By inspecting the slopes of the log-log curve (i.e. $\partial \log P_l(t) / \partial \log t$) at 20 minutes (i.e. short lifetime) and 600 minutes (i.e. long lifetime), it is also found that the shape of our model matches to those curves obtained by other researchers based on the measurements in 2001, 2002 and 2004. Therefore it is observed that the session lifetime distribution of Gnutella might be an invariant characteristic independent of the technological advancement or protocol change during the period 2002 to 2005.

Keywords: Gnutella, peer-to-peer, session lifetime distribution

1 Introduction

Understanding the underlying structure of a network is not a new problem. Since the discovery that a number of natural networks have the small-world phenomena [17], i.e. the node degree distribution follows power law, a number of measurement studies on various different kinds of networks have been conducted. The Faloutsos brothers have found that our Internet is a small world network [5]. Albert, Jeong and Barabasi [1] found that the World Wide Web is a small world network. Since small world networks are claimed to be natural

networks, a number of measurement studies regarding the network topology have thus been done in the past few years.

Gnutella is a peer-to-peer overlay network for file sharing. In terms of number of users, Gnutella is the 3rd largest peer-to-peer network after eDonkey and FastTrack¹. In spite of an interest in the power law topology, Jovanovic did a crawl on a subset of Gnutella peers and counted their node degrees. In [6], he reported that the node degree distribution of Gnutella is basically similar to a power law distribution. He then attempted to re-construct the network topology of Gnutella based on the neighbor relationships amongst the peers' IPs being crawled. As Gnutella is a dynamic network, it is hardly to validate how matches the Gnutella at any instance. Ripeanu *et al.* [9], by comparing their measurement statistics obtained in November 2000 and May 2001, found that the node degree distribution of Gnutella was changing over the period of measures. It changed from a power law distribution to a kind of mixture between power law and Poisson. Since Gnutella protocol v6.0 has been released, Sum *et al* [12] did a crawl in 2003 and found that the node degree distribution of the ultrapeer-ultrapeer subnet follows a two power-laws distribution.

Owing to its popularity and the consumption large amounts of network resources, various measurement studies have been carried out in the last few years aiming at improving the file sharing capabilities as well as better monitoring the network traffic generated by the peers [3, 10, 11].

Saroiu *et al* [10] and Sen & Wang [11] independently did measurements in 2001 capturing various characteristics of Gnutella. Using a comprehensive understanding on the extra traffic being generated by the Napster and Gnutella, their numbers of file shared, upload/download statistics and how often a peer connect and disconnected from the system, Saroiu *et al* did an active measure by periodically (in a 7 minutes interval) crawling each system to collect snapshots. They found that there is significant heterogeneity and lack of cooperation across peers participating in these systems. Ignoring the sessions with uptime less than 7 minutes², they have found that most sessions are of short uptime and the median session duration is approximately 60 minutes. Also as many as 25% of peers in which has no file shared (the so-called free riders). They simply connect to the Gnutella download files and then disconnect.

On the other hand, Sen & Wang [11] took a passive approach by analyzing 800 million *flow*³ records collected at multiple routers across the ISP's network for Direct-Connect, FastTrack and Gnutella overlay networks over a

¹It is based on the August 31, 2005 statistics in Slyck.com.

²It corresponds to the un-measurable portion.

³In accordance with Sen & Wang, a flow is defined to be an unidirectional sequence of packets between a particular source and destination IP address pair.

period of 3 months. They have observed less than 10% of the IPs contributes around 99% of the total traffic volume. Also, many traffic characteristics such as traffic volume and connection time per unique source/destination IP pair are extremely skewed. Only a small fraction of peers are persistent over long time periods. As no actual figures or tables of results related to Gnutella have been presented, it is hard to cross validate the claim they have made with the results obtained in [10].

Chu *et al* [3] presented a study (based on active measurement conducted over 3 months in 2002) of the popularity of files stored and transferred among peers in Napster and Gnutella and indicated that the popularity of files is skewed. Furthermore, they have also found that distribution of session uptime follows a log-quadratic function. This observation has also been found in a recent measurement studied by Stutzbach & Rejaie [14] who did an active probing measurement for a large number of IPs with 7 minutes probing cycle on October 2004 over a period of 15 hours.

Although other measurement results have been reported, such as using a fast probing system [15, 16] that can quickly capture millions of IPs within 7 minutes and on the bootstrapping time behavior [2, 7], not much work has been done on the change of the session lifetime distributions over a long period of time (in terms of years). Session lifetime is of paramount importance for the understanding of the topological evolution of Gnutella, such as the node degree distribution and the network connectivity. As mentioned in [15], all these factors will as a result affect the overlay structure, the resiliency of the overlay and the simulation of such a network. Eventually, it will affect the selection of design parameters for a P2P network. In this regards, the primary objective of this paper is to present a measurement result we did in 2003 on the session lifetime of ultrapeers and give a comparative discussion on this parametric model with the models obtained by Chu *et al* [3] and by Stutzbach & Rejaie [14]. In this paper, session lifetime and session uptime are used interchangeably. They correspond to the duration a peer is connecting to Gnutella. This is also called node availability in some other papers.

The remainder of the paper will be organized as follows. In the next section, we will give an overview on the connection mechanism of Gnutella. Then in Section 3, the methodology of measurement and results on node degree distribution and session lifetime distribution will be elucidated⁴. The estimated model is presented in Section 4. A discussion of the results compared with others will be presented in Section 5. The conclusion of the paper will be presented in Section 6.

⁴Throughout the paper, lifetime and session lifetime are used interchangeably.




Figure 1. Peer connection.

2 Peer Connection Mechanism (Bootstrapping)

In accordance with Gnutella protocol [4], while a new peer (Peer A) is trying to make connection to the existing Gnutella network, Peer A can either send a XTRY or PING message to an online Gnutella peer (Peer B) asking for the IP addresses of its neighbors, Figure 1. If XTRY is used, Peer B returns with a list of its direct neighbors. If PING is used, Peer B will broadcast the request to its direct neighbors. Its direct neighbors will then broadcast the request to their direct neighbors again and again until the message has been propagated TTL (time-to-live) hops away.

All their IP addresses will then be propagated backward to Peer B based on the PONG protocol. After receiving these PONG messages, Peer B will again send a PONG message to Peer A. Whether XTRY or PING-PONG protocol is used, Peer A can select a number of IP addresses from the list and make connections.

There are several remarks for this connection mechanism. The first remark is about the *seed IP* — where a peer can get the online peers' IP addresses. Once a peer, say Peer A, wants to connect to Gnutella, it needs to send a message to an online Gnutella peer. Basically, the software developer will set two default locations for a peer to find online peers' IP addresses. The first location is the local cache, which stores all the IP addresses a peer previously connected to. The second location is a well-known global server like limewire.com. It stores the IP addresses of all the current online peer nodes. Whenever a peer wants to connect, it can get the IP addresses from these locations and select one IP address as the seed for XTRY or PING.

The second is about *the number of seed nodes*. Suppose there are two available seed nodes, say X and Y, found in the local cache. In the current setting,

Peer A is unlikely to ask Peer Y for further information if Peer X has been asked and replied with a list of IP addresses. Therefore, we can assume only one seed node will be XTRYPed whenever a peer node wants to make a connection to Gnutella. This assumption has also been strengthened by a recent paper on simulated Gnutella [13]. Whenever a peer connects to 2 or more seed nodes, the shape of the node degree distribution in log-log scale should be single modal.

The third remark is about *the connectedness of Gnutella*. If the network is not connected, with a number of isolated components, the availability of information from other components will definitely be hidden from Peer B. Obviously, Gnutella is a highly dynamic network. Peer nodes can go offline whenever they finish downloading files, or encounter software or hardware failure. When a large number of peer nodes go offline, the connectedness will no longer be guaranteed. Therefore, once the network has been disconnected, asking information from only one seed node cannot help in recovering the network connectedness.

3 Measurement Method and Results

The measurement was carried out in August 2003. The measurements of the node degree and the node lifetime are carried out in two consecutive phases.

3.1 *Ultrapeers collection*

First, a list of IP addresses (IPs) is obtained from limewire.com. Starting from this list of IPs, we send out two messages: *Gnutella Ultrapeer: True*. Once a *Gnutella 200 OK* message has been received, we immediately send out a *PING message* to crawl its neighbor IPs. Then a *Gnutella 200 OK* message is sent to these neighbor IPs. We repeat this process (crawl) until some 100000 IP addresses have been collected as a master list.

3.2 *Handshaking*

As a Gnutella peer can define itself as either a blocked ultrapeer or non-blocked ultrapeer, node degree measurement is a bit complicated. To facilitate fast capturing, 20 threads are out simultaneously for handshaking. As the IP addresses collected in the last step might be outdated in the database, our program only threads out to those IP addresses that have no more than 60 minutes of log time.

For each thread, (i) an IP address is randomly selected from the database and a bootstrapping message with *Ultrapeer: True* is sent to it. (ii) If the

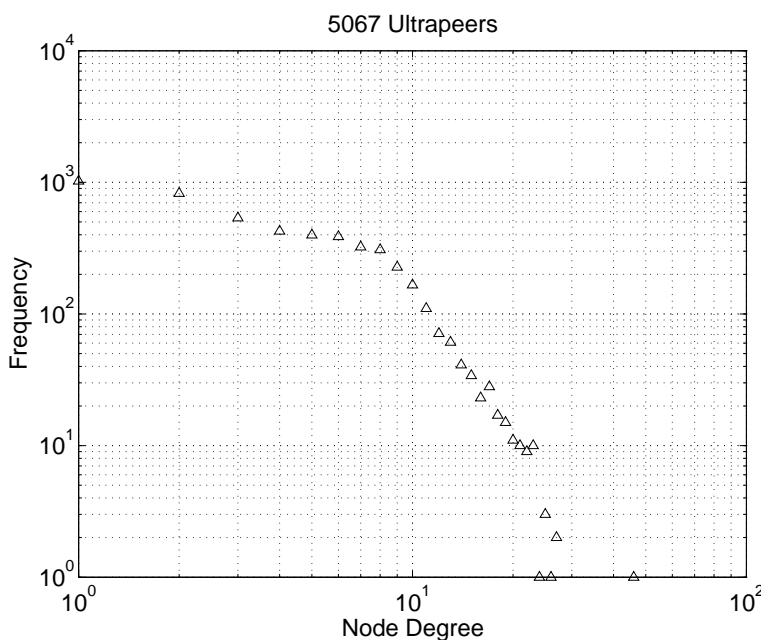


Figure 2. Node degree distribution of the ultrapeer-ultrapeer subnet of Gnutella.

node sends back a message with *Gnutella 200 OK*, it confirms that this is a successful connection. Otherwise, we assume that the node is either inactive or acting as a block node. (iii) If the node is non-blocked, a message PING TTL 2 is sent to that node asking for neighbor IPs. (iv) If the node is a blocked node, a new IP is selected randomly from the remaining IP addresses and the same message *Ultrapeer: True* is sent to it. The received PONG replies are stored in a database.

By analysis on the data obtained, we can then identify whether a node is a leaf, a blocked ultrapeer, or a non-blocked ultrapeer. To maintain uniformity, our study only focuses on the ultrapeer-ultrapeer node degree distribution and eventually 5087 non-blocked ultrapeer nodes are captured and identified. Their node degree distribution is shown in Figure 2. Observed from the curve, the distribution looks like a mixture of two power-law distributions :

$$P_n(k) \propto \begin{cases} k^{-0.8} & k \leq 5 \\ k^{-4} & k > 5 \end{cases} \quad (1)$$

where k is the node degree and $P_n(k)$ is the probability of a node with degree k .




Figure 3. Typical cases that are measured.

3.3 Online session time measurement

The methodology of collecting session time information is similar to the way we explore online peers. 7500 sample nodes are randomly selected from the database. Another database of 7500 records with $8 \times 24 \times 4$ fields in each record.

Every 15 minutes these IPs are tested by sending bootstrap messages to them. If a node replies with *Gnutella 200 OK* or *Gnutella 503*, it is defined as active and a '1' will be saved in the corresponding record and the corresponding time field. If there is no reply from the node a '0' will be written. Figure 3 shows a few typical cases that are captured from the measurement. Let IP_1 , IP_2 and IP_3 be the IP addresses to be measured. In IP_1 , two sessions *A* and *B* are captured. In IP_2 , three sessions have been captured during the measurement duration. As the session measured in the beginning of measurement has already been alive long before the measurement started, it is excluded from our statistics accounting for the lifetime distribution. Similarly, the session that is still alive when the measurement has been ended will be excluded from our statistics. Therefore, only session *C* will be considered as a valid candidate. Similarly, only session *D* and *E* will be considered for lifetime distribution. After excluding all the invalid sessions, there is another problem that we need to decide. Should sessions *A* and *B* be considered as a single session or two? To give an answer to this problem, we take the following approach.

After the program has ran for eight days (768 hours), 7500 records similar to the following line were obtained.

```
110011111101111000000000111111110000.....
```

The first '1' corresponds to the first measurement. Since we have no information about when it starts, we ignore the first two ticks. In the 3rd and 4th ticks, there is no reply. We assume that during that period of time, the IP has not been assigned with an active Gnutella peer. So, the 5th tick will be the starting uptime of a peer. But there is one question left: When is the end of a




Figure 4. Number of sessions counted against the number of no-replies.

session?

Like the above example, $\dots 0011111101111000\dots$, we assume '1 0 0' indicates the end of a session. Whenever an IP does not reply for 2 or more consecutive ticks (i.e. ≥ 30 minutes), it is assumed to be offline. Later, if the IP is found to be active again, it is treated as a new peer.

So, in this example, we assume the IPs have been occupied two times by two different Gnutella peers. Their uptimes are 11×15 minutes and 8×15 minutes respectively. Another reason why we measure at every 15 minutes is because we assume that active uptime of an ultrapeer should be long, in order of hours.

To confirm our assumption, the number of sessions counted for the following four cases has been shown in Figure 4.

- (i) Assume the session is off if no reply has been received for 1 time, i.e. whenever '1 0' has encountered.
- (ii) Assume the session is off if no reply has been received for 2 times, i.e. whenever '1 0 0' has encountered.
- (iii) Assume the session is off if no reply has been received for 3 times, i.e. whenever '1 0 0 0' has encountered.
- (iv) Assume the session is off if no reply has been received for 4 times, i.e. whenever '1 0 0 0 0' has encountered.

Using '1 0' as an end of a session, the short lifetime sessions will be over-counted.




Figure 5. Lifetime distribution directly measured from Gnutella for 768 hours. 1 NR: A node is assumed offline if no reply has been received for one trial. 2 NR: A node is assumed offline if no reply has been received for two consecutive trials. 3 NR: A node is assumed offline if no reply has been received for three consecutive trials. 4 NR: A node is assumed offline if no reply has been received for four consecutive trials.

The normalized lifetime distribution is shown in Figure 5 and Figure 6 shows the corresponding cumulative distribution. The distribution obtained for case 1 is a bit different. For case 1, more 'short uptime' data was counted. Eventually, it makes the slope of the 'log-log' plot in the 'short uptime' region steeper. By observing the number of sessions being counted against the numbers of no-replies in Figure 4, we assume that the number of sessions in the case '1 NR' is over-counted.

4 Estimated model

Therefore, the lifetime data corresponding to the case '2 NR' is used to model the lifetime distribution for ultrapeers. By inspecting the shape of the cumulative lifetime distribution in Figure 6, a reasonable model for the distribution

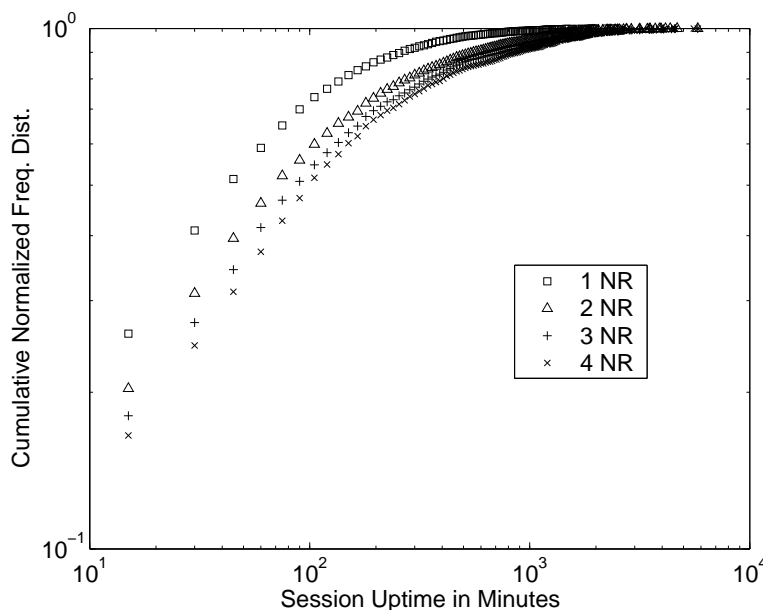


Figure 6. Cumulative lifetime distribution directly measured from Gnutella for 768 hours. 1 NR: A node is assumed off-line if no reply has been received for one trial. 2 NR: A node is assumed off-line if no reply has been received for two consecutive trials. 3 NR: A node is assumed off-line if no reply has been received for three consecutive trials. 4 NR: A node is assumed off-line if no reply has been received for four consecutive trials.

is power law with exponential cut off, defined as follows :

$$P_l(t; \alpha, \beta) \propto \left(\frac{t}{\beta}\right)^{-\alpha} \exp(-t/\beta). \quad (2)$$

To estimate the parameters of α and β , one cannot directly use least squares fit to the log-log curve of

$$\frac{P_l(kT; \alpha, \beta)}{\sum_k P_l(kT; \alpha, \beta)}$$

and the data in Figure 4. Here T is the time interval between two consecutive measures. To be exact, we define normalized distribution $\bar{P}_l(kT)$ as follows :

$$\bar{P}_l(kT) = \frac{\int_0^T \int_0^T P(t_e|kT + t_s; \alpha, \beta) dt_e P(t_s) dt_s}{\sum_{k \geq 1} \int_0^T \int_0^T P(t_e|kT + ts) dt_e P(t_s) dt_s}, \quad (3)$$

Here t_s and t_e correspond to the unknown session starting offset and the session ending offset, Figure 7. $P(t_e|kT + t_s; \alpha, \beta)$ is the conditional probability that




Figure 7. Definition of a session.

the session will be alive for t_e extra time when it has already been alive for $kT + t_s$. $P(t_s)$ is probability that a new session starts at t_s before its first measurement. We assume that the arrival of a new peer is uniformly random. That is to say, $P(t_s) = 1/T$. For $P(t_e|kT + t_s)$, there is no simple close form that we can use. Therefore, the double integral in Equation (3) is calculated numerically by the following equation :

$$\sum_{i=1}^m \sum_{j=1}^m P_l(kT + i\delta + j\delta; \alpha, \beta),$$

where $m\delta = T$. By Law of Large Number, it can be proved that

$$\sum_{i=1}^m \sum_{j=1}^m P_l(kT + i\delta + j\delta; \alpha, \beta)(\delta)^2 \approx \int_0^T \int_0^T P(t_e|kT + t_s; \alpha, \beta) dt_e P(t_s) dt_s,$$

for small δ and for all $k \geq 1$. Thus, the corresponding cumulative distribution, denoted by $\bar{C}_l(kT; \alpha, \beta)$, can be given by

$$\bar{C}_l(kT; \alpha, \beta) = \frac{\sum_{g=1}^k \sum_{i=1}^m \sum_{j=1}^m P_l(gT + i\delta + j\delta; \alpha, \beta)}{\sum_{g \geq 1} \sum_{i=1}^m \sum_{j=1}^m P_l(gT + i\delta + j\delta; \alpha, \beta)}. \quad (4)$$

The best estimate of (α_0, β_0) will be defined as

$$(\alpha_0, \beta_0) = \arg \min_{(\alpha, \beta)} \left\{ \sum_{k \geq 1} (\log \bar{C}_l(kT; \alpha, \beta) - \log C_l(kT))^2 \right\}, \quad (5)$$

where $\log C_l$ s are the values shown in Figure 6.

To avoid costly computation on the gradient of $\log \bar{C}_l$, we set $m = 150$ and




Figure 8. The normalized probability distribution of the estimated model versus the actual data measured. The solid line corresponds to the estimated model while the dots correspond the 2 NR data.

exhaustive search for⁵

$$\alpha \in 1.00, 1.01, 1.02, \dots, 1.50;$$

$$\beta \in 100, 110, 120, \dots, 990.$$

Eventually, it is found that the best estimate for (α, β) is $(1.12, 780)$. In other word, the session lifetime distribution is given by :

$$P_l(t; 1.12, 780) \propto \left(\frac{t}{780}\right)^{-1.12} \exp\left(-\left(\frac{t}{780}\right)\right). \quad (6)$$

Figure 8 compares the curves of $\bar{P}_l(kT)$ and the normalized distribution with 2 NR. The comparison on their cumulative counterparts is shown in Figure 9.

It should be noted that the time scale in Equation (6) is in minute. One can see that the value β is equal to 13 hours. That means, most Gnutella peers are only alive for no more than 13 hours.

⁵The same results are obtained even when $m = 300$.




Figure 9. The cumulative distribution of the estimated model versus the actual data measured. The solid line corresponds to the estimated model while the dots correspond the 2 NR data.

5 Related works since 2001

As mentioned earlier in the paper, many measurement studies on Gnutella have been done in the last couple of years, with particular interests in node degree distribution, traffic volume, file distributions and node uptime. Only a few of them provided information about session lifetime. Table 2 summarizes some of these works since 2001.

Saroiu *et al* [10], one of the earliest group of researchers who based their work on a create-based approach⁶ captured the session uptimes of Gnutella nodes. In their measurement, it is found that 50% of the peers never remain online for more than one hour and the median session duration is approximately 60 minutes. As they claimed, it is similar to the time a user takes to download a small number of music files.

On the other hand, Sen & Wang [11] took a passive approach by collecting the *flow* information from Cisco's *Netflow* systems. For Fast Track, it is found that about 50% sessions are of lifetime less than 10 minutes. Although they have claimed that the result for Gnutella is similar, no actual figure has been reported. A drawback of their analysis is that the lifetime of each session is

⁶The captured traces are divided into two halves. The reported durations are for the sessions that started in the first half, and finished in either the first or second half.

Works	MM/YY	T (mins)	Node size	Duration
Saroiu <i>et al</i> [10]	05/01	7	17125	2.5 days
Sen & Wang [11]	09/01	0.02	–	1 day
Chu <i>et al</i> [3]	05/02	10	5000	3 days
This paper	08/03	15	7500	8 days
Cramer <i>et al</i> [2]	xx/03*	10	665723#	5 weeks
Wong [18]	11/03	10	7500 [†]	8 days
Stutzbach & Rejaie [14]	04/04 – 10/04	2, 7	1.3 M	72, 31 hours
Stutzbach & Rejaie [15]	10/04 – 12/04	7	1.3 M	2 days

* In accordance with the year the paper has been included in CiteSeer.

On Gwebcache nodes only.

[†] Since the set of probing nodes are initially crawled from Gwebcache, large proportion of the set is Gwebcache nodes.

Table 1. Summary of measurement works in between 2001 – 2004, with session lifetime information.

over-counted. As *NetFlow* collects only *Transport Layer* information, it is not possible to identify whether a *Flow* is generated after or during bootstrap. The connection duration (Figure 8a in [11]) reported is essentially the subtotal of bootstrapping duration and session uptime.

In accordance with their measurement conducted in 2002, Chu *et al* [3] is almost the first research group to suggest that session lifetime of a Gnutella peer is of log quadratic form. With reference to a measurement done in 2004, Stutzbach & Rejaie [14] re-confirm this claim. To accurately capture snapshots of Gnutella, Stutzbach & Rejaie developed a system called Cruiser which can capture a large number of peers within just a few minutes. By probing to 1.3 millions peers for 72 hours with bin size 2 minutes⁷, they have identified that the session lifetime distribution of the top-level peers follows a power law distribution with exponent -1.756 .

In another probing measure for 31 hours with bin size 7 minutes, they have shown the session lifetime distribution of all peers can be characterized by either a 2-power-law distribution or a log-quadratic distribution. In [15], they have further illustrated that the shapes of the session lifetime distribution captured at 5 different periods of time, from October to December of 2004, are not exactly the same. The exponents of the session lifetime in the range of [12 minutes, 1440 minutes] are between -1.30 to -1.68 .

While studying the bootstrapping characteristics, Cramer *et al* did a measure on the Gwebcache nodes [2] some time in 2003 to identify their actual availability. We reproduce the session lifetime distribution they obtained in Figure 10 for reference. The straight lines (not from their original figure) are

⁷Bin size corresponds to the time interval between two consecutive probe.

added to illustrate the slopes of the tangents at the points $t = 20$ minutes, 600 minutes and 2000 minutes respectively. Wong did a similar measure on the lifetime of nodes that were crawled from Gwebcache and its neighbors in the August of 2003 [18]. Both results show that the lifetime distribution is like a power law distribution with exponential cutoff and the exponent for lifetime in between $t = 200$ minutes, 2000 minutes is around -2 . For the sake of clarification, the results obtained in the previous studies are summarized in Table 2 and Table 3.

Works	$p(t)$	$\partial \log p(t) / \partial \log t$
Chu <i>et al</i> [3]	$t^{-0.61} \exp(-0.07(\log t)^2)$	$-0.61 - 0.14 \log t$
This paper	$t^{-1.12} \exp(-t/780)$	$-1.12 - t/780$
Wong [18]	$t^{-1.998}$	-1.998
Stutzbach & Rejaie [14]	$t^{-1.756}$	-1.756
Stutzbach & Rejaie [14]	$t^{-0.126} \exp(-0.149(\log t)^2)$	$-0.126 - 0.3 \log t$

All logs are natural logarithms.

Table 2. Summary of session lifetime models.

Table 3 lists the exponents

$$\frac{\partial \log p(t)}{\partial \log t}$$

of the models at $t = 20$ minutes and $t = 600$ minutes. For Cramer *et al*'s data, the values are estimated by inspecting Figure 10. One can observe that the

Works	MM/YY	$T(\text{mins})$	20 mins	600 mins
Chu <i>et al</i> [3]	03/02 - 05/02	10	1.03	1.50
This paper	08/03	15	1.15	1.90
Cramer <i>e al</i> [2]	xx/03	10	1.45 [†]	2.0 [†]
Wong [18]	11/03	10	-	2.0
Stutzbach & Rejaie [14]	09/04	2	1.756	1.756
Stutzbach & Rejaie [14]	10/04	7	1.18	1.90
Stutzbach & Rejaie [15]	10/04	7	1.30, 1.34	1.30, 1.34
Stutzbach & Rejaie [15]	11/04	7	1.45	1.45
Stutzbach & Rejaie [15]	12/04	7	1.60, 1.68	1.60, 1.68

[†] The data is estimated by inspecting the curve provided by Cramer *et al* in [2], see also Figure 10 in this paper.

Table 3. The exponents locally at t equals to 20 and 600 minutes.




Figure 10. Gwebcache uptime distribution adapted from Cramer *et al* [2]. Slope of LINE 1, LINE 2 and LINE 3 are approximately -1.445, -2.0 and -4 respectively.

model estimated in this paper (based on our measurement in August 2003) is rather consistent with the one obtained by Stutzbach & Rejaie in [14] (based on their measurement in October 2004).

Figure 11 shows the scatter plot of the exponents with respect to the date of the measurement. One can see that the exponents are within $[1, 2]$ for all $20 \text{ mins} \leq t \leq 600 \text{ mins}$. Except the results obtained by Chu *et al* in 2002 and by Cramer *et al* in 2003 specifically for Gwebcache nodes, the remaining data fit well within the range. In accordance with Chu *et al*'s results, the slope of the distribution is not that deep, while the slope in Cramer *et al*'s distribution is deeper than others. We do not have a definite answer for this case but we believe it is due to the protocol difference and the slow Internet access rate at that time. For the latter case, we suspect that it is due to the node sampling problem.

Excluding the two special cases measured in 2002 and 2003 respectively, one could make a statement refined from [14]. Stutzbach & Rejaie have suggested that *the power law distribution of session times is independent of the application and the purpose of application. Rather it is determined by application-independent aspects of user-behavior.* Here, we add that *user-behavior in general is a time invariant characteristic.*

Before proceeding to the conclusion of the paper, a few notes should be commented about the mathematical models for the lifetime distribution.

- Measurement of very short lifetime session is notoriously difficult. Hence, the modeling of distribution head is still an open problem. Without such




Figure 11. Exponents against the date the measurement conducted. Here January 2002 is set to be the 1st month. The '+' dots correspond to the values at $t = 20$ minutes while the triangular dots correspond to the values at $t = 600$ minutes.

information, all measurement results reported so far cannot truly reflect the actual distribution, whether in terms of probability density (i.e. PDF) or cumulative distribution (CDF).

- One major obstacle in using either 2-power-laws distribution, log-quadratic distribution or power law with exponential cut off is that they are not integrable, i.e.

$$\int_{t_0}^{\infty} P_l(t) dt$$

is undefined when $t_0 \rightarrow 0$. This leads to the conclusion that the analysis on the growth and the size of Gnutella is impossible, unless a less fitting model like Gamma distribution [12] is employed.

6 Conclusion

In this paper, a measurement result obtained in 2003 regarding the session lifetime of Gnutella peers has been presented. The measurement is based on the active probing method. The results on node degree distribution and session lifetime distribution have been shown. We found that the session life-

time statistics can fit to a power law distribution with exponential cutoff, $P_l(t) \propto (t/780)^{-1.12} \exp(-t/780)$ for all session lifetime $t \geq 15$ minutes. The result is similar to the model estimated by Stutzbach & Rejaie in a measurement done in 2004. To study the evolutionary change in session lifetime, a survey on work done in between year 2001 and 2004 on session lifetime distribution has been elucidated. Comparing our 2003 measurement result with the others in the literature, it is also found that session lifetime seems to be a time invariant characteristic. For sessions with relatively long lifetime (around 600 minutes), the exponent of the power law fit is around -1.9 to -2.0. For sessions with relatively short lifetime (around 20 minutes), the exponent is around 1.2. In view of the recent findings about session lifetime, we make the following observation.

After 2002, when Gnutella v0.6 has been released and broadband access has become the common Internet access medium, the session lifetime seems to be an application independent property and it looks like an invariant property of Gnutella, which is depended solely on the user behavior. And users behavior should also be time invariant.

One difficulty in simulating the topological formation of Gnutella due to the non-integrable property of power law or power law with exponential cutoff has also been remarked. It brings out that searching for a better distribution model should be a worthwhile direction for further investigation. Finally, one should take a note on the measurement results reported by Cramer *et al* [2] and Wong [18] on Gwebcache lifetime distribution. In their measurement, it is found that the lifetime of Gwebcache is relatively short (the exponent is -2.0) compared with Gnutella peers in general. As Gwebcache is a technique to facilitate peer discovery, its unavailability can affect the performance of the overall setup. Further investigation along this line seems to be inevitable. Moreover, whether the results apply to other networks in the same category should also be explored in the future.

Acknowledgement

The authors would like to express their gratitude to the reviewers for their valuable comments.

References

- [1] R. Albert, H. Jeong and A.-L. Barabasi, Diameter of the World Wide Web, *Nature*, 401:130-131, 1999.
- [2] C. Cramer, K. Kutzner, and T. Fuhrmann, Bootstrapping Locality-Aware P2P Networks *Proceedings of the IEEE International Conference on Networks (ICON 2004)*, vol. 1, pp. 357-361, 2004.
- [3] J. Chu, K. Labonte and B.N. Levin, Availability and locality measurements of peer-to-peer file systems, *IT COM: Scalability and Traffic Control in IP Networks II Conferences*, July 2002.

- [4] Clip2 Distributed Search Services, The Gnutella protocol specification v0.4, <http://dss.clip2.com>.
- [5] M. Faloutsos, P. Faloutsos and C. Faloutsos, On power-law relationships of the Internet topology, *Computer Communication Review* 29:251, 1999.
- [6] M.A. Jovanovic, Modeling Large-scale Peer-to-Peer Networks and a Case Study of Gnutella, Master Thesis Department of Electrical and Computer Engineering and Computer Science, University of Cincinnati, April 2001.
- [7] P. Karbhari, M. Ammar, A. Dhamdhare, H. Raj, G. Riley and E. Zegura, Bootstrapping in Gnutella: A Measurement Study, *Proceedings of the Passive and Active Measurements Workshop*, Antibes Juan-les-Pins, France, April 2004.
- [8] R. Matei, A. Iamnitchi, P. Foster, Mapping the Gnutella network, *IEEE Internet Computing*, Vol. 6(1)50 - 57, 2002.
- [9] M. Ripeanu *et al.*, Mapping the Gnutella Network, Jan/Feb 50-57, *IEEE Internet Computing*, 2002.
- [10] S. Saroiu, P.K. Gummadi, S.D. Gribble, A measurement study of Peer-to-Peer file sharing systems, *Proceedings of the Multimedia Computing and Networking (MMCN)*, San Jose, January, 2002.
- [11] S. Sen and J. Wang, Analyzing peer-to-peer traffic across large networks, *IEEE/ACM Transactions on Networking*, Vol.12(2), 219-232, 2004.
- [12] J. Sum *et al.*, A Study of the Connectedness of Gnutella, PDPTA'04, Las Vegas, 2004.
- [13] J. Sum and S.C. Lau, A novel connection mechanism for P2P, PDPTA'04, Las Vegas, 2004.
- [14] D. Stutzbach and R. Rejaie, Towards a better understanding of churn in peer-to-peer networks, Technical Report CIS-TR-04-06, Department of Computer Science, University of Oregon, November 2004. Available at the URL www.barsoom.org/~agthorr/blog/Publications.shtml.
- [15] D. Stutzbach and R. Rejaie, Characterizing Churn in Peer-to-Peer Network, to appear at *Internet Measurement Conference'06* October 25-27, Brazil, 2006. Available at the URL www.barsoom.org/~agthorr/blog/Publications.shtml.
- [16] D. Stutzbach and R. Rejaie, Characterizing the two-tier Gnutella topology, *SIGMETRIC'05*, 2005.
- [17] D. Watts, *Small Worlds*, Princeton University Press, 1999.
- [18] M.F. Wong, Statistical Analysis of worm spreading and resilience study in P2P, Project Report, Department of Computing Hong Kong Polytechnic University, 2004.