

ID Repetition in Kad

Jie Yu^{1*}, Chengfang Fang², Jia Xu², Ee-Chien Chang², Zhoujun Li³

¹Department of Computer Science, National University of Defense Technology, China

²Department of Computer Science, National University of Singapore, Singapore

³School of Computer Science and Engineering, Beihang University, China
yj@nudt.edu.cn, {c.fang,xujia,change}@comp.nus.edu.sg, lizj@buaa.edu.cn

Abstract

ID uniqueness is essential in DHT-based systems as peer lookup and resource searching rely on ID-matching. Many previous works and measurements on Kad do not take into account that IDs among peers may not be unique. We observe that a significant portion of peers, 19.5% of the peers in routing tables and 4.5% of the active peers (those who respond to Kad protocol), do not have unique IDs. These repetitions would mislead the measurements of Kad network. We further observe that there are a large number of peers that frequently change their UDP ports, and there are a few IDs that repeat for a large number of times and all peers with these IDs do not respond to Kad protocol. We analyze the effects of ID repetitions under simplified settings and find that ID repetition degrades Kad's performance on publishing and searching, but has insignificant effect on lookup process. These measurement and analysis are useful in determining the sources of repetitions and are also useful in finding suitable parameters for publishing and searching.

1 Introduction

In DHT-based Peer-to-Peer (P2P) systems, each node or object (e.g. a keyword or file) is assigned with an identity (ID), which plays a crucial role during the lookup, publishing and searching processes. In some implementations of DHT, mechanisms are incorporated to ensure that the identifiers assigned to peers are unique. Many research works [2, 3, 4, 5, 18] have been devoted to this, including centralized certification and distributed certification. However, these certifications either require certain private information of users,

such as cellular phone number [18] or email account [2], which are not acceptable for many P2P anonymous users, or need to combine with network layer information [5], e.g. IP address, sacrificing node mobility. For example, in Pastry, each identifier is derived from the IP address or the public key [12]. On the other hand, some implementations choose not to enforce ID uniqueness. One example is Kad, which is an implementation of Kademlia protocol [11] and has millions simultaneous users as to date. Each peer in Kad is free to choose an ID of 128-bit string as its identifier. The intention of such design decision is probably to reduce startup time and to support node mobility [22]. However, without enforcing uniqueness, multiple peers can share the same ID. This leads to the interesting question of whether such sharing and repetition of ID appears in actual DHT-based systems.

Indeed, we observe that, in Kad, there is a significant amount of repetition — 19.5% of the peers in routing tables and 4.5% of the *active peers* (i.e., peers who respond to the BOOTSTRAP requests.) do not have unique IDs. There are a few possible causes of the repetition, including intentional attacks on selected keywords, bugs in client software, crawlers, sensors or botnets that are setup to gather data, etc. There are many studies and analysis performed on Kad, including peer behaviors and distribution [17, 19], publishing and searching [16], lookup [20], security [18, 22], routing table [24], etc. Most of these works do not consider the repetition of the 128-bit ID. Without taking such repetition into account, measurements on Kad might be distorted, and parameters chosen for certain algorithms, for example publishing and searching processes, may not perform as expected.

To measure and study the effects of the ID repetition, we first gather ID related information of peers in Kad using a crawler for several weeks. Our observations and analysis can be summarized as below:

1. Among the peers who appeared in all routing ta-

*Work done during internship in National University of Singapore.

bles, a low fraction (about 50%) of them respond to Kad messages. We believe that such low fraction of active peers is mainly due to long lifetime of contacts in routing tables (a peer left while its outdated contact is still in some routing tables), and UDP port aliasing (a peer frequently switches its UDP ports), instead of IP address aliasing or peers located behind NAT or firewall.

2. There is a significant amount of repetition — 19.5% of the peers appeared in routing tables and 4.5% of the active peers do not have unique IDs. Interestingly, the repetitions in routing tables follow a Zipf-like distribution. Such ID repetition would impact some measurements in previous works, such as lifetime or geographic distribution.
3. There are some *silent* groups of peers in the routing tables. Each group contains a large number (more than 100) of peers which all share the same ID, and none of them are active peers. Possibly, it is due to some modified clients or malicious tools connected to Kad that do not follow the standard Kad protocol.
4. We analyzed the effect of ID repetition on lookup and searching under some simplified and reasonable assumptions. The analysis shows that lookup performance will not be degraded while the searching performance of certain targeted keywords will be degraded significantly, and these two results are verified empirically. The analysis also shows that by having slightly more peers to store the keywords, the lookup process can tolerate much more ID repetitions. Such analysis is useful in understanding Kad and provides a guide in choosing appropriate parameters for the relevant algorithms.

2 Background and Related Work

Kad is the first DHT implementation and it has millions simultaneous users as to date. Recently, Kad has been suggested to be the underlying infrastructure of other large-scale applications to increase the scalability and decrease the deployment cost, such as Second Life [21]. In Kad, the distance between two IDs is defined by the “XOR metric”, which is a weighted Hamming distance. Each peer keeps a routing table which contains a list of contacts, and each contact mainly consists of 3 components: an ID, the respective address and the contact type. The contact type is an integer in $\{0, 1, 2, 3, 4\}$, and represents the level of availability. Value 0 means the best availability and value 4 means the worst. The lifetime of a contact with type value 4,

3, 2, 1 or 0 is 0 minute, 2 minutes, 1 hour, 1.5 hours or 2 hours respectively.

Given a *target*, say k , *lookup* is carried out to locate the peer in Kad network whose ID is closest to k . Kad employs an iterative process to lookup the target by querying routing tables in neighboring peers. Stutzbach *et al.* [20] gave a detailed description and performance analysis on the lookup process. To *publish* a keyword, a 128-bit string k is derived from the keyword and treated as an ID. Information of the keyword is to be published and stored in γ peers whose IDs are closest to k , where γ is a parameter and is usually set to 10. To find these peers, typically, the iterative process of lookup is carried out with k as the target. Among the list of peers visited during lookup, the keyword is to be published to γ closest peers who are willing to keep the information. To *search* for a keyword, similar process is carried out to determine γ closest peers. The search is successful if at least one of these γ peers is located.

Many works have focused on the performance and security of Kad. K. Kutzner and T. Fuhrmann [9] measured the *IP address aliasing* in Overnet because of the use of dynamic IP addresses. During a two-week measurement, they found some IDs appeared to associate with more than 100 different addresses. M. Steiner *et al.* [17] observed *Kad ID aliasing* where a client changes its ID after one or several sessions. They found that most of this phenomenon appeared in China. We investigated this issues and believe that this is mostly due to the new versions of two popular download tools in China: FlashGet and Thunder. Both tools have integrated Kad but do not save ID information locally and randomly generate it in every session. To the best of our knowledge, there is no detailed investigation of ID repetition in Kad in the literature. Perhaps the most related work is the observation given by M. Steiner *et al.* [17, 15], that the distribution of ID over the 256 8-bit zones is not uniform as expected in Kad. They suggested that this could be due to software bugs.

Partially due to the lack of a certification service to verify peers’ identities [2], Kad is vulnerable to a large spectrum of attacks, such as sybil attack [6], eclipse attack [14], routing table and index poisoning [10, 25], DDoS attack [22, 25], etc. Sybil attack [6] introduces multiple malicious peers into a distributed system, aiming to take control of the whole overlay network. M. Steiner *et al.* [18] discussed that sybil attack in Kad can be exploited to spy on publishing and searching traffic, eclipse contents or perform DDoS attacks. Eclipse attack [14] aims to separate a set of victim nodes from the rest of the overlay network. ID repetition discussed in this paper can be considered as a spacial case of

eclipse attack, while it just tries to isolate the searching of keywords or files. P. Wang *et al.* [22] proposed to hijack routing table of clients using spoofed contacts to perform DoS attack to Kad network. To counter with vulnerabilities in Kademia protocol, L. M. Aiello *et al.* [2] proposed *Likir*, a framework that is built on top of Kademia and includes an identity based scheme and a secure communication protocol. It may provide an effective defense against above attacks, but is not easy to implement.

3 Measurement

3.1 Data Gathering

To obtain a snapshot of Kad, we deployed a crawler similar to *Blizzard* [17, 19]. The main difference is that we use BOOTSTRAP request instead of ROUTING request to traverse Kad network. Using BOOTSTRAP could be more effective, since a single BOOTSTRAP response contains 20 contacts while one ROUTING response contains at most 11 contacts [16]. We record two sets of peers: one set contains the peers who appear in routing tables; the other set contains the *active peers*, i.e., the peers who respond to the BOOTSTRAP requests. These two sets are analyzed separately to investigate ID repetitions in routing tables and among active peers. Note that most previous works [17, 19, 15, 20, 24] on measurement of Kad are done on the first set (i.e. all peers in routing table). To distinguish different peers, we keep track of the IP address and UDP port for each peer found. We also record TCP port and Kad version of each peer for a more in-depth analysis. The snapshots were obtained from 10 Feb 2009 to 15 Apr 2009, everyday around 18:00 GMT (daytime in East Asia) and 06:00 GMT (nighttime in East Asia). For each snapshot, we look for IDs, each of which maps to multiple (IP address, UDP port) tuples, and classify the corresponding peers as peers with repeated ID.

Each snapshot we collected contains about 2.5 to 3.8 million peers in routing tables and about 1.2 to 1.8 million active peers. This is slightly less than the numbers measured by M. Steiner *et al.* from March 2007 to May 2008, which are 3 to 4.5 million and 1.5 to 2 million respectively [17, 19]. The size of Kad during 06:00 GMT, corresponding to nighttime in East Asia, is about 20% to 30% larger than the size during the corresponding daytime. This is consistent with the measurements done by M. Steiner *et al.* [17, 19].

Note that there are only about half of peers in routing tables responding to the BOOTSTRAP requests.

(a) Repetition among peers in routing tables

Number of repetitions	Mar 24 Tue	Mar 25 Wed	Mar 26 Thu	Average fraction
1	2689181	2696062	2691355	81.0%
2	335060	336282	341612	10.1%
3	91179	89637	92409	2.7%
4	46772	44968	45164	1.4%
5	26705	25955	26365	0.8%
6	15942	15720	16158	0.5%
7	12117	10815	10850	0.4%
8	9216	8664	9096	0.3%
9	7947	7173	7713	0.2%
10	6450	6560	6910	0.2%
11 – 20	31979	29990	29656	0.9%
21 – 50	15827	14718	15408	0.5%
51 – 100	4310	4565	4384	0.1%
101 – 1000	14088	12939	14338	0.4%
> 1000	13887	14524	22358	1.2%

(b) Repetition among active peers

Number of repetitions	Mar 24 Tue	Mar 25 Wed	Mar 26 Thu	Average fraction
1	1497743	1463648	1473345	95.5%
2	26850	25786	28588	1.7%
3	5064	4869	4833	0.3%
4	4180	4208	4268	0.3%
5	3770	3825	3815	0.3%
6	3312	3042	3558	0.2%
7	2646	2296	2632	0.2%
8	2568	2032	1976	0.1%
9	1683	1332	1593	0.1%
10	1330	900	1220	0.1%
11 – 20	4052	2678	2769	0.2%
21 – 50	1472	1356	1305	0.1%
51 – 100	432	653	867	0.1%
101 – 1000	6429	6096	6273	0.4%
> 1000	7591	6929	7464	0.4%

Table 1. Distribution of peers by number of repetitions

This low proportion may be due to the following reasons:

C1: The routing table of each peer keeps every contact for a period of time. When a node left, its information may still remain in routing tables until its lifetime expires. When we send a message to such node, there would be no response.

C2: Peers located behind NAT or firewalls can not receive any request messages directly and hence are unable to respond.

C3: It takes 25 to 40 minutes for our crawler to collect information of all peers. During this period, some peers may change their IP addresses and thus can not receive request messages. This is possible as the Internet service providers may be running DHCP [9, 19] and it is known as *IP address aliasing*.

C4: Peers can selectively or completely ignore Kad

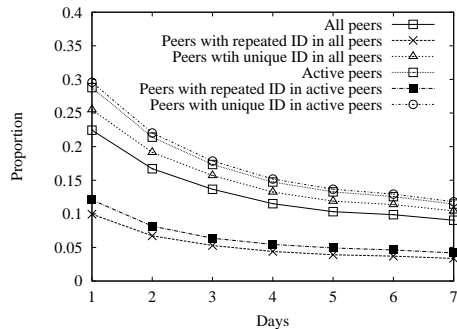


Figure 1. Proportions of peers that also appeared in the reference day. Here “all peers” refers to peers in routing tables.

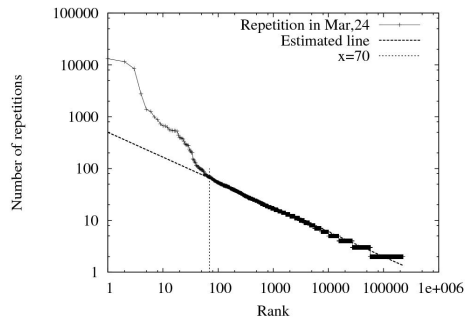
messages. These peers could be “selfish” ones who make use of Kad network but contribute little or none, or “bad” ones who exploit Kad for other purposes.

M. Steiner *et al.* [17] suggested that this low active rate of peers in routing tables is mainly due to *C2*. However, in Section 4.1 we will show that *C2* contributes a little to the fraction of unresponsive peers, while *C1* and *C4* are the main reasons. Furthermore, we observe a special case of *C4*, which we call *UDP port aliasing*.

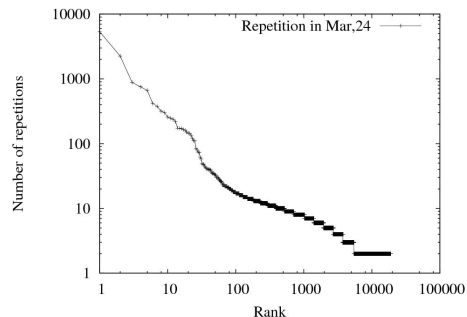
3.2 Distribution of Repetitions

The fraction of peers without unique ID is significant, about 19.0% among peers in routing tables and 4.5% among active peers in every snapshot we obtained. In Section 4, we will analyze why these two numbers are so different from each other. If there are m peers having k as their ID, let us say that *the number of repetitions of k is m* . Table 1 shows the distribution of peers by the number of repetitions of their IDs among peers in routing tables and active peers respectively. The first row shows the number of peers with unique ID, and the second row shows the number of peers whose IDs repeated twice, and so on. The percentages in the last column are the average over 3 days.

Note that in Table 1 the distributions are similar over a few days. We next look at the lifetime of individual peer. We first choose a snapshot (Mar 24, 2009) as reference. For each of the next following seven days, the fraction of peers that also appear in the reference snapshot are recorded. Fig.1 shows the fractions among peers with repeated ID and unique ID for both sets of peers in routing tables and active peers respectively.



(a) Repetition among peers in routing tables



(b) Repetition among active peers

Figure 2. Number of repetitions per ID in log-log scale.

Clearly, there is a gap between the two curves for each set. Previous works on lifetime [17, 19] have not taken ID repetition into consideration. It would be interesting to further investigate the lifetime of different types of peers.

The number of repetitions of individual ID, sorted in decreasing order, are shown in Fig.2 in the log-log scale. It is interesting that Fig.2(a) shows a Zipf-like distribution (i.e. the probability of a randomly chosen ID having rank i is proportional to $\frac{1}{i^a}$ where a is a constant) of repetitions among peers in routing tables, for ranking 70 onwards. Note that the graph is in log-log scale, and the fraction of peers having ID with rank at most 70 is actually small, less than 1.7%. Hence, majority of the peers follows the Zipf-like distribution, with $a \approx 0.48$. Shamma *et al.* [13] observed the Google document frequency of the terms formed a Zipf-like distribution and M. Steiner *et al.* found that the distribution of keywords stored in Kad also follows Zipf-like distribution [16]. However, from Fig.2(b), the repetitions among active peers do not follow a Zipf-like distribution strictly. It is not clear why these two distributions are different and why the first distribution fits nicely to the Zipf-distribution.

(a) Most popular IDs in routing tables

Rank	ID (in hex format)	No. of Re- petitions	No. of Countries	Country with Lar- gest No. of Peers	Rank among Active Peers	No. of Repetitions in Active Peers
1	09262ce48db41838ce94c80cdaab3fab	13175	27	CHN (96%)	–	0
2	00000000000000000000000000000000	11557	87	CHN (39%)	1	5345
3	ab3d5a03c4892c603dd9beda87eda8d8	8492	48	CHN (95%)	2	2246
4	02ac8fc8a3e4caba1b1b520a623d5732	2751	23	CHN (95%)	4	752
5	5ad0327057fd75e85bf687011af12c3c	1379	5	ISR (98%)	3	879
6	e188d20e843abb978cd6eb24c591b846	1265	10	ISR (96%)	921	7
7	80f9ade2d68dc455efca6364cb9e9a31	984	17	FRA (95%)	5	664
8	dc173f913d2b41156fcb22373c4dcb74	875	2	CHN (99%)	13	230
9	efbef3a7f4406f876ee4dde0077ce8d0	726	12	KOR (96%)	11	247
10	9c2d76dd27f26ff1fe4f2a6544d53582	667	10	KOR (94%)	12	238

(b) Most popular IDs among active peers

Rank	ID (in hex format)	No. of Re- petitions	No. of Countries	Country with Lar- gest No. of Peers	Rank among All Peers	No. of Repetitions in All Peers
1	00000000000000000000000000000000	5345	66	CHN (29%)	2	11557
2	ab3d5a03c4892c603dd9beda87eda8d8	2246	26	CHN (96%)	3	8492
3	5ad0327057fd75e85bf687011af12c3c	879	5	ISR (98%)	5	1379
4	02ac8fc8a3e4caba1b1b520a623d5732	752	14	CHN (96%)	4	2751
5	80f9ade2d68dc455efca6364cb9e9a31	664	14	FRA (94%)	7	984
6	f111cbe0f1f0efcc2ceb37b27b0a6fb5	419	11	ARG (52%)	11	662
7	ad445207f1062287ef54341e0c110d7a	376	14	ITA (88%)	12	629
8	75b7002f8f0c5f0e8124db7e2b79ae0e	317	7	ESP (94%)	15	541
9	25ee26b837958af0307a086a51420368	301	28	ITA (28%)	13	559
10	672b3949cc995b6c47865eb796436380	257	11	ESP (86%)	14	551

Table 2. Ten most popular IDs in routing tables and among active peers. Here “all peers” refers to peers in routing tables.

Table 2 lists the details of the first few IDs with the largest number of repetitions among peers in routing tables and among active peers respectively. We investigate the groups of peers with the same ID. Here we use the file (config/ip-to-country.csv) provided in eMule client [1] to resolve IP addresses to countries. Note that for a single ID, the corresponding peers could spread over multiple countries. For some IDs, the corresponding peers are mainly concentrated in a single country, for example, the first and third rows in Table 2(a). On the other hand, some IDs are widely spread, for example, the sixth and ninth rows in Table 2(b). We do not observe any interesting patterns in the geographic distributions of various IDs. Note that none of the 13,175 peers in the group listed in the first row of Table 2(a) are active. We call such group a *silent group* and will further investigate it in Section 4.3.

Fig.3 shows the geographic distribution of peers in the top 8 countries. The distribution of peers in each subset (i.e. all peers, peers with repeated ID in all peers, peers with unique ID in all peers, and so on. See Fig.3.) is different from others to a certain degree. Previous works on geographic distribution of Kad [17, 19] have not taken this into consideration. Especially, we find that among active peers, the percentage of peers

in Spain is about 16.5%, while it is 35.4% among peers with repeated IDs, twice more than that of the global percentage.

3.3 Behaviors of Different Peers

To compare the behaviors of active peers with and without unique ID, we randomly select 1000 peers with ID repeated at least 3 times and 1000 peers with unique ID from the set of active peers, and then send routing requests of a randomly chosen ID and publishing requests of a random keyword. For those who agreed to store the keyword, we later send a request to retrieve the keyword. The result is shown in Table 3. The reason that we receive fewer routing responses is that peers will not respond if they do not have relevant routing information. From Table 3, we can see that most active peers with repeated ID carry out Kad protocol as honestly as the peers with unique ID. This is a useful observation which leads to an assumption we made in Section 5: the majority of the peers with repeated ID would carry out routing, publishing and searching honestly.

We also search for these published keywords using eMule [1] (version 0.49b), which is a standard Kad

Table 3. Fraction of successful operations

	Routing	Publishing	Retrieving	Searching By eMule
Unique	25.1%	89.8%	99.1%	56.7%
Repeated	27.2%	83.1%	97.8%	13.2%

client software. Different from the process described in the previous paragraph, whereby a keyword is retrieved from a known peer, an eMule client, given a keyword, first carries out the lookup process to locate the relevant peers, and then carries the retrieval process. As shown in the last column of Table 3, the success rate of hitting keywords stored in peers with repeated ID is significantly lower than that in peers with unique ID. In Section 5.2 we will formulate the success rate of searching in the situation with repeated IDs.

4 Analysis of repetition

In this section, we try to answer two questions: (1) Why up to about 50% of peers in routing tables do not respond to the BOOTSTRAP request messages? (2) Why the fraction of repetitions among peers in routing tables (i.e. 19.5%) is much larger than that among active peers (i.e. 4.5%)?

The first question has been overlooked in previous works. However, it is important to study it because understanding it helps to design more effective management policies of Kad’s routing tables, so as to increase the active rate of contacts. It also helps in solving the second question. The analysis to the second question provides insightful understanding to different kinds of repetitions.

4.1 Analysis of the low active rate of peers in routing tables

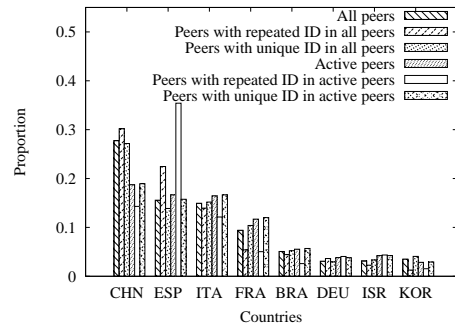
As discussed in Subsection 3.1, there are four possible causes, i.e., $C1$, $C2$, $C3$ and $C4$, that induce the low rate of active peers in routing tables. In the next few subsections, we will estimate the fraction of peers affected by each of the four causes.

4.1.1 Estimate the fraction of departure peers ($C1$)

To measure the fraction of contacts with different types (i.e. types range from 0 to 4) in routing table, we add a module into the eMule software to record the routing table every 10 minutes. From a two-week monitoring, the average fraction of each type is shown

Table 4. Fraction of each type of contacts in a routing table

Type 0	Type 1	Type 2	Type 3	Type 4
82.3%	7.7%	7.1%	2.3%	0.6%

**Figure 3. Histogram of geographic distribution of different peers. Here “all peers” refers to peers in routing tables.**

in Table 4. It shows that most contacts in routing tables are of type 0, which means that the lifetime of most contacts in routing table are 2 hours. M. Steiner *et al.* [19] monitored a subset of live peers in Kad and obtained an approximated CDF (cumulative distribution function) of session times. We can then estimate the fraction of departure peers as: $\frac{1}{2} \sum_{i=0}^4 [\text{fraction of type } i \times \text{CDF}(\text{lifetime of type } i)]$. As a result, the fraction of departure peers is estimated as 24.2%.

4.1.2 Estimate the fraction of blocked peers ($C2$)

One of the design goals of contact type is to distinguish peers located behind NAT or firewall, since these peers do not participate in routing or storing published information. Normally, these peers have contact types with values not smaller than 3. So, we can estimate an upper bounded of the fraction of these blocked peers in routing tables by the fraction of peers with contact type values 3 or 4, which add up to 2.9%.

4.1.3 Estimate the fraction of peers with IP address aliasing ($C3$)

In this situation, each peer appears as sharing the same ID with some others since its ID, as well as UDP port, TCP port and version do not change with IP address. Firstly, we argue that this fraction should be insignifi-

cantly small: (i) the “lease time” of the DHCP is typically longer than an hour, for instance, the default value for Windows Server 2003 is 8 days and most ISPs use the value of 24 hours, while it takes only 25 to 40 minutes for our crawler to collect information of all peers. (ii) even if the IP address of a peer is changed, there is a delay of a few minutes for this information to propagate to other routing tables.

Secondly, to support our argument, we estimate this fraction by investigating the distributions of peers in routing tables and active peers. Consequently, we estimate that the fraction of peers with IP address aliasing is smaller than 2.0%. We leave the details to our full paper [23].

4.1.4 Estimate the fraction of intentional unresponsive peers (C4)

Inactive peers due to the causes of $C1$, $C2$ and $C3$ can be treated as peers who do not intend to remain inactive, but can not respond due to the ISP or are wrongly classified due to outdated information in the routing tables. We are not aware of other major causes and treat the remaining inactive peers not belong to $C1$, $C2$ and $C3$ as peers who are programmed not to respond to the BOOTSTRAP messages. Hence, we estimate the fraction of such peers to be $50\% - (24.2\% + 2.9\% + 2.0\%) = 20.9\%$. These peers might selectively or completely ignore Kad messages.

In sum, the low rate of active peers in routing tables is mainly due to the long lifetime of contacts in routing tables and intentional silence of some peers in response to Kad messages, instead of IP address aliasing or peers located behind NAT or firewall.

4.2 UDP port aliasing

A Peer in Kad communicate with others through UDP port, which should remain the same across all sessions. However, we observe that among peers in routing tables, there are some groups, such that all peers in each group have the same ID, IP address, TCP port and version, but different UDP ports. We call this phenomenon as *UDP port aliasing*, in contrast to IP address aliasing [9] and ID aliasing [17]. Such phenomenon is due to the frequent UDP port switching by some peers. One possible motivation of such behavior is that those peers try to avoid being connected by other peers in Kad. The fraction of peers with UDP port aliasing among peers in routing tables is 9.2%. Most of these peers contribute to low active rate discussed in the previous section.

Note that UDP port aliasing is not due to different peers located behind the same NAT and thus show the same public IP address in Kad. The reasons are: (i) different peers behind the same NAT still have different IDs; (ii) we have verified that at most one peer in each group responds to our requests at some point in time; and (iii) we have estimated that the fraction of blocked peers in routing table is about 2.9%, which is much smaller than 9.2%.

4.3 Silent groups

There are many groups of peers where all peers in a group share the same ID and they all do not respond to the BOOTSTRAP request message. For further investigation, routing and publishing messages are sent to them and yet there are still no response. These peers appear to be “silent” in Kad network. The fraction of these peers accounts to 8.7% among peers in routing tables. These silent groups can be divided into the following two types: (i) groups with peers that do not intend to remain silent and the repetition is due to “artifacts” of network protocols. This includes the multiple peers recorded due to IP address aliasing, which contribute to 1.0 % as discussed in Section 4.1.3. We are not aware of other mechanisms and thus take IP address aliasing as the main cause for groups in this type. (ii) groups with peers that intentionally do not respond to standard Kad requests. The second type of groups can be further divided into two subtypes: groups with UDP port aliasing and groups with peers across multiple ISPs. We further calculate that the first subtype contributes to 3.8% and thus the second subtype contributes the rest 3.9%.

We focus on the biggest silent group with ID corresponding to the first row in Table 2(a). Since 96% of these peers are from China, we further investigate its distribution and find that they are widely distributed over most provinces of China. We can see from Fig.4 that its group size varies between 12,000 and 15,000, and the size at weekend is about 15% larger than that at weekday. We first randomly pick a snapshot of this group as reference. For each of the next seven days, the number of peers that share the same IP addresses or the same (IP addresses, UDP port) tuples and also appear in the reference snapshot are recorded. Fig.4 shows that the number of peers sharing the same IP addresses is about 4,000 while the number of peers sharing the same (IP addresses, UDP port) is only about 600, which means that these peers are possibly still involved in UDP port aliasing.

We suspect that most of these silent groups corresponds to some modified clients or malicious tools con-

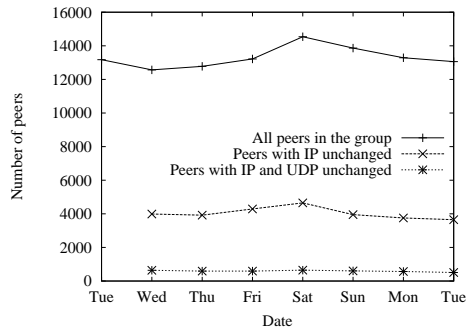


Figure 4. The number of peers of the biggest silent group over a week.

nected to Kad, which do not follow the standard Kad protocol. They just try to take advantage of Kad network and do not want to accept connections from normal peers. It is possible that they publish their peer information into routing tables of Kad, so that other members in the same silent group can locate them easily. Storm Worm [7, 8] is a well-known peer-to-peer botnet worm. It propagates via spam and communicates using overnet protocol, which is also based on the Kademlia protocol and is very similar to Kad. Since overnet has been dead for more than two years, it is possible that the attackers have switched or are trying to switch from overnet to Kad since it just requires a little modification on the original Storm Worm. One of our future works is to investigate their relationship.

4.4 Genuine Repetition

As analyzed above, peers can change their IP addresses or UDP ports dynamically before their information registered at routing tables expires. These peers are treated as multiple peers sharing the same ID by our crawler. This repetition is “false repetition” and it contributes to less than 11.2% among peers in routing tables. The others are different peers possibly across multiple ISPs while sharing the same ID, which is called *genuine repetition*. It contributes to more than 8.3% among peers in routing tables. Note that this fraction is still larger than 4.5% measured in active peers. It is because (i) some silent groups never respond to our BOOTSTRAP requests and (ii) the total number of active peers is only half of all peers in routing tables. This genuine repetition may be due to intentional attacks on selected keywords, bugs in client software, crawlers, sensors or botnets that are setup to gather data, etc.

Note that although peers with “false repetition” do

obey the ID assignment in Kad, they would still make lookup and searching less effective since Kad will ignore the other peers with the same ID, once it has found one. However, we can avoid this by changing Kad protocol slightly, such that it validates the aliveness of each contact before using it. In next section we will analyze the effect of genuine ID repetition on lookup and searching.

5 Effect of ID Repetition

In this section, we investigate the effect of genuine repetition. We assume that peers with repeated ID carry out Kad’s protocol honestly. Thus, its information will appear in other routing tables, and it will respond to lookup, searching and publishing processes. This assumption is supported by the observation made in Section 3.3. Here we first investigate the effect on lookup, followed by the effect on publishing/searching. Let n be the number of peers who are online at a particular time, and m the number of distinct IDs.

5.1 Lookup

D. Stutzbach *et al* [20] gave the expected number of hops¹ required during lookup as follows:

$$1 + \frac{\log_2 n - 7.33}{6.58} \quad (1)$$

When taken ID repetition into consideration, the expected number of hops in (1) may be slightly overestimated: a more accurate estimation should use m instead. ID repetition improve the look up performance because: (i) information of repeated peers appear more often in routing tables comparing to a normal single peer, and (ii) repeated peers lives longer in Kad. These certainly will affect the performance of lookup, and are worthy of further investigation.

The total number of distinct IDs in snapshots we gathered is about 86% of the number of peers, which is between 2.5 to 3.8 million. Applying the new formula, the expected number of hops is between 3.08 and 3.21. To verify this, we modify the eMule client to perform lookup of 10 randomly chosen targets every 10 minutes. This experiment is run for a week and the average number of hops is 3.14, which is within estimation. Actually, replacing n by $m = 0.86n$ in a logarithm has a very small effect. In sum, ID repetition does not degrade the lookup performance.

¹Note that the number of hops is not proportional to the time taken or number of peers visited. Nevertheless, it serves as a good indicator of the performance of lookup.

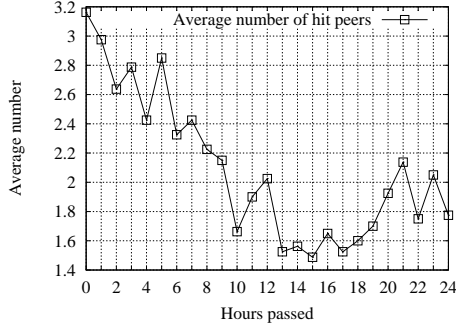


Figure 5. The average number of hit peers by the eMule client over time.

5.2 Publishing and Searching

Now we investigate the process of publishing and searching under ID repetition. To publish a keyword or object with hash k , its information is to be stored in γ peers whose IDs are closet to k , and to search for it, information is to be retrieved from γ peers whose IDs are closet to k . There are two groups of peers of interests, the group P_p of γ peers involved in the publishing and another group P_s of peers involved in the searching. The search fails if there is no common peers in these two groups, i.e., $P_p \cap P_s = \emptyset$. Typically, in Kad, γ is chosen to be 10. To measure the number of peers hit in publishing/searching in present Kad, we modify the eMule client to publish a key to peers with unique ID and then search for these keys every hour for 24 hours² starting from 19:00 GMT, Mar 24, 2009. For each search, the number of peers holding the key is recorded, and 100 experiments are conducted in parallel. Fig.4 shows the average number of hit peers. Since as long as there is a single hit, the search is successful. Hence, the search/publishing is rather robust on average. However, as shown in Section 3.3, the success rate of hitting keywords stored in peers with repeated ID is significantly lower than that in peers with unique ID.

Let us now consider the scenario where IDs are repeated. When a keyword k is successfully published once to the ID i , only one peer with ID i will store this keyword. Therefore, the chance of successful retrieval from this ID is $\frac{1}{f(i)}$, where $f(i)$ is the number of repetitions for ID i . This could be exploited by adversary to prevent some specific files to be shared in Kad. The best strategy to “isolating” a keyword is discussed in our full paper [23], and the effect of such isolation is as shown in table 5.

²The experiment was conducted for 24 hours since peers typically clear the stored keywords every 24 hours.

Table 5. Success probability when 10 peers are holding the keyword

Peers known by adversary	4	6	8	10
Attacked by 20 peers	77.1%	69.3%	60.5%	49.8%
Attacked by 40 peers	75.6%	65.1%	51.2%	33.5%

Table 6. Success probability when 20 peers are holding the keyword

Peers known by adversary	8	12	16	20
Attacked by 80 peers	94.1%	87.8%	76.1%	55.8%
Attacked by 160 peers	93.6%	85.8%	69.5%	36.2%

In practise, the adversary may not find all the peers that store the keyword. The success probability will increase when there are more peers storing the keyword that are not found by the adversary. Table 5 gives an illustration of this where there are a total of 10 peers holding the keyword, all peers have a 20% probability to be visited during a search, and the adversary can generate 20 or 40 peers and always choose the best strategy. From table 5, we can see that when the adversary’s information is limited, the additional resource for generating more peers does not help much.

To summarize, ID repetition will reduce the performance of Kad by increasing possibility of a failure in searching for a published key. The probability of such failure depends on the number of repetitions, the number of IDs holding this key, and the probability that those IDs are visited during a search.

6 Mitigation

As mentioned in the introduction, it is possible to enforce ID uniqueness, for example, by requiring that each ID is derived from the IP address and UDP port using a cryptographic-secure hash function. However, employing this mechanism will lose node mobility since a new ID must be generated after the IP address is changed. Furthermore, it is tedious to support nodes located behind NAT, since these nodes need to obtain the translated IP address and port number to compute the ID.

A straightforward improvement is to treat the tuple (ID, IP address, UDP port) as the real identifier during publishing and searching. However, we have conducted a simulation with the ID repetition distribution according to the snapshot in Mar 24, 2009 and

find that doing this actually *reduces* the probability of successful search.

One effective way to mitigate the ID repetition is to recruit more peers to store the keywords during publishing. Table 6 shows the probability of successful search for a key held by 20 peers with 20% probability to be visited during a search and attacked by an adversary with 80 peers and 160 peers. Comparing to Table 5, although the adversary has quadruple resources while the keyword is published only in twice more peers, the success probability increases. In general, from the approximation form of (3), to maintain the same probability, the adversary need to generate peers to increase both $f(i_j)$ and d . In other words, the resource required by the adversary is “quadratic” in the resource required by the publisher.

7 Conclusion and Future Work

Kad does not enforce ID uniqueness and we observed that a significant fraction of peers have their IDs shared with others. Fortunately, the design of Kad is robust enough that even with the presence of ID repetition, there is no significant degradation in lookup performance. Although the performance of searching and publishing is severely degraded for targeted keywords, our analysis shows that by publishing the keyword to slightly more peers, much more repetitions are required to degrade the searching performance. Nevertheless, if the mobility of peers is not a requirement, it would be desirable to prevent the problem by enforcing the ID uniqueness in the design. Although our studies classify the peers based on their behaviors, it is still not clear what are the sources of repetitions. It is interesting to further investigate and identify the sources. It is also interesting to find out whether such repetition exists in other DHT-based systems.

References

- [1] Emule. <http://sourceforge.net/projects/emule/>.
- [2] L. Aiello, M. Milanese, G. Ruffo, and R. Schifanella. Tempering Kademia with a Robust Identity Based System. In *Proceedings of P2P'08*, 2008.
- [3] K. Butler, S. Ryu, P. Traynor, and P. McDaniel. Leveraging Identity-Based Cryptography for Node ID Assignment in Structured P2P Systems. *IEEE Transactions on Parallel and Distributed Systems*, 2008.
- [4] M. Castro, P. Druschel, A. Ganesh, A. Rowstron, and D. Wallach. Secure routing for structured peer-to-peer overlay networks. In *Proceedings of OSDI'02*, 2002.
- [5] J. Dinger and H. Hartenstein. Defending the sybil attack in p2p networks: Taxonomy, challenges, and a proposal for self-registration. In *Proceedings of ARES'06*, 2006.
- [6] J. Douceur. The sybil attack. In *Proceedings of IPTPS'02*, 2002.
- [7] T. Holz, M. Steiner, F. Dahl, E. Biersack, and F. Freiling. Measurements and mitigation of peer-to-peer-based botnets: a case study on storm worm. In *Proceedings of the LEET'08*, 2008.
- [8] C. Kanich, C. Kreibich, K. Levchenko, B. Enright, G. Voelker, V. Paxson, and S. Savage. Spamalytics: an empirical analysis of spam marketing conversion. In *Proceedings of ACM CCS'08*, 2008.
- [9] K. Kutzner and T. Fuhrmann. Measuring large overlay networks—the overnet example. In *Proceedings of KiVS'05*, 2005.
- [10] J. Liang, N. Naoumov, and K. Ross. The index poisoning attack in P2P file-sharing systems. In *Proceedings of INFOCOM'06*, 2006.
- [11] P. Maymounkov and D. Mazieres. Kademia: A peer-to-peer information system based on the xor metric. In *Proceedings of IPTPS'02*, 2002.
- [12] S. Rhea, D. Geels, T. Roscoe, and J. Kubiatowicz. Handling churn in a DHT. In *USENIX Annual Tech. Conf.*, 2004.
- [13] D. Shamma, S. Owsley, K. Hammond, S. Bradshaw, and J. Budzik. Network arts: exposing cultural reality. In *Proceedings of WWW'04*, 2004.
- [14] A. Singh, T. Ngan, P. Druschel, and D. Wallach. Eclipse attacks on overlay networks: Threats and defenses. In *Proceedings of INFOCOM'06*, 2006.
- [15] M. Steiner, E. Biersack, and T. Ennajary. Actively monitoring peers in KAD. In *Proceedings of IPTPS'07*, 2007.
- [16] M. Steiner, W. Effelsberg, T. En-Najjary, and E. Biersack. Load reduction in the kad peer-to-peer system. In *Proceedings of DBISP2P'07*, 2007.
- [17] M. Steiner, T. En-Najjary, and E. Biersack. A global view of kad. In *Proceedings of IMC'07*, 2007.
- [18] M. Steiner, T. En-Najjary, and E. Biersack. Exploiting KAD: possible uses and misuses. *Computer Communication Review*, 2007.
- [19] M. Steiner, T. En-Najjary, and E. Biersack. Long Term Study of Peer Behavior in the KAD DHT. *IEEE/ACM Transactions on Networking*, 2009.
- [20] D. Stutzbach and R. Rejaie. Improving lookup performance over a widely-deployed DHT. In *Proceedings of INFOCOM'06*, 2006.
- [21] M. Varvello, E. Biersack, and C. Diot. P2P Second Life: experimental validation using Kad. In *Proceedings of INFOCOM'09*, 2009.
- [22] P. Wang, J. Tyra, E. Chan-Tin, T. Malchow, D. Kune, N. Hopper, and Y. Kim. Attacking the Kad network. In *Proceedings of SecureComm'08*, 2008.
- [23] J. Yu, C. Fang, J. Xu, E. Chang, and Z. Li. ID repetition in Kad. Available at <http://www.comp.nus.edu.sg/~changecc/publications/P2P09TR.pdf>.
- [24] J. Yu and Z. Li. Active Measurement of Routing Table in Kad. In *Proceedings of DAS-P2P'09*, 2009.
- [25] J. Yu, Z. Li, and X. Chen. Misusing Kademia Protocol to Perform DDoS Attacks. In *ISPA'08*, 2008.