

Comparing paedophile activity in different P2P systems

Raphaël Fournier, Thibault Cholez, Matthieu Latapy, Clémence Magnien
Isabelle Chrisment, Ivan Daniloff, Olivier Festor

June 6, 2012

Abstract

Peer-to-peer (P2P) systems are widely used to exchange content over the Internet. Knowledge on paedophile activity in such networks remains limited while it has important social consequences. Moreover, though there are different P2P systems in use, previous academic works on this topic focused on one system at a time and their results are not directly comparable.

We design a methodology for comparing *KAD* and *eDonkey*, two P2P systems among the most prominent ones and with different anonymity levels. We monitor two *eDonkey* servers and the *KAD* network during several days and record hundreds of thousands of keyword-based queries. We detect paedophile-related queries with a previously validated tool and we propose, for the first time, a large-scale comparison of paedophile activity in two different P2P systems. We conclude that there are significantly fewer paedophile queries in *KAD* than in *eDonkey* (approximately 0.09% vs 0.25%).

1 Introduction

Paedophile activity is a crucial social issue and is often claimed to be prevalent in peer-to-peer (P2P) file-sharing systems [3, 7]. However, current knowledge on paedophile activity in these networks remains very limited. Recently, research works have been conducted to improve this situation by quantifying paedophile activity in *Gnutella* and *eDonkey*, two of the main P2P systems currently deployed [5, 6]. They respectively conclude that 1.6% and 0.25% of queries are of paedophile nature, but these numbers are not directly comparable as the authors use very different definitions and methods. Such

comparisons are of high interest though, since differences in features of P2P systems, such as the level of anonymity they provide, may influence their appeal for paedophile users.

In this paper, we perform for the first time such a comparison. We focus on the *KAD* and *eDonkey* P2P systems, which are both widely used and differ significantly in their architecture: while *eDonkey* relies on a few servers, *KAD* is fully distributed. This lack of centralisation may lead users to assume that *KAD* provides a much higher level of anonymity than *eDonkey*. Comparing the two systems sheds light on the influence of a distributed architecture on paedophile behavior and increases general knowledge on paedophile activity in P2P systems.

Section 2 describes our datasets and how we collected them. Section 3 presents our comparison of the amount of paedophile queries in *KAD* and *eDonkey*. Section 4 focuses on an important feature of paedophile activity: ages entered in queries. Finally, in Section 5 we infer the fraction of paedophile queries in *KAD* from the one in *eDonkey*, and Section 6 presents our conclusions.

2 Datasets and measurements

In order to compare paedophile activity in two different P2P systems, we first need appropriate datasets, the collection of which is a challenge in itself. In *KAD* and *eDonkey*, different kinds of measurements are possible.

In *eDonkey*, servers index files and providers for these files, and users submit keyword-based queries to servers to seek files of interest to them [4]. By monitoring such a server, one may collect all those queries [1]. Here, we record all queries received by two of the largest *eDonkey* servers during a three-month period in 2010. The servers are located in different countries (France and Ukraine) and have different filtering policies: the French server indexes only non-copyrighted material, while the Ukrainian server openly indexes all submitted files. Monitoring two such different servers will allow us to compare them in order to know if server policy impacts our results.

To collect *KAD* data, we use the HAMACK monitoring architecture [2], which makes it possible to record the queries related to a given keyword by inserting distributed probes close to the keyword ID onto the *KAD* distributed hash table. We supervise 72 keywords, which we choose to span well the variety of search requests entered in the system, with a focus on paedophile activity: a set of 19 *paedophile* keywords (*babyj*, *babyshivid*, *childlover*, *childporn*, *hussyfan*, *kidzilla*, *kingpass*, *mafiasex*, *pedo*, *pedofilia*, *pedofilo*, *pedoland*, *pedophile*, *pthc*, *ptsc*, *qqaazz*, *raygold*, *yamad*, *youngvideomodels*), which are

known to be directly and unambiguously related to paedophile activity in P2P networks; a set of 23 *mixed* keywords (*1yo, 2yo, 3yo, 4yo, 5yo, 6yo, 7yo, 8yo, 9yo, 10yo, 11yo, 12yo, 13yo, 14yo, 15yo, 16yo, boy, girl, mom, preteen, rape, sex, webcam*) frequently used in paedophile queries but also in other contexts (for instance, *Nyo* stands for *N years old* and is used by both paedophile users and parents seeking games for children of this age); and a set of 30 *not paedophile* keywords (*avi, black, christina, christmas, day, doing, dvdrip, early, flowers, grosse, hot, house, housewives, live, love, madonna, man, new, nokia, pokemon, rar, remix, rock, saison, smallville, soundtrack, virtual, vista, windows, world*) used as a test group and *a priori* rarely used in paedophile queries. The sets of keywords result of the work on paedophile query detection presented in [6]. Notice that our set of keywords contains mainly common English words (*love, early, flowers*), but some are in other languages (*saison, pedofilia*), and some are also brand names (*pokemon, nokia*).

Because of the differences in architectures of the two networks and of the measurement methodologies, we obtained very different datasets, which are not directly comparable: in *eDonkey*, we observe all queries from a subset of users whereas in *KAD* we only observe queries related to a given keyword, but from all users. In addition, based on various versions of *KAD* clients, the measurement tool only records the queries containing a monitored keyword placed in first position or being the longest in the query. As a consequence, with a short keyword such as *avi*, a name extension for video files, we almost only record queries in which it is the unique keyword, because otherwise it most likely is neither the longest nor the first word in any query. In order to obtain comparable datasets, we therefore limit our study to a subset of our datasets: the queries composed of exactly one word among the 72 keywords we monitor.

As a result of this construction process, we obtain three datasets, which we call *edonkeyFR*, *edonkeyUA* and *KAD*. They contain 241,152, 166,154 and 250,000 queries respectively, all consisting of a unique keyword from our list of 72 monitored keywords, which ensures that they are comparable. The server corresponding to the *edonkeyFR* dataset is located in France, while the one corresponding to *edonkeyUA* is in Ukraine. Notice moreover that their large sizes make us confident in the reliability of our statistical results presented hereafter.

3 Amount of paedophile queries in *eDonkey* versus *KAD*

The most straightforward way to compare the paedophile activity in different systems certainly is to compare the fraction of paedophile queries in each system. Figure 1 presents the fraction of queries for each category of keywords. This plot clearly shows that there are very distinct search behaviors in the two networks, since values obtained for the *paedophile* and *not paedophile* categories significantly differ between *KAD* and the two *eDonkey* datasets. More surprisingly, the fraction of paedophile queries is significantly lower in *KAD* than in *eDonkey* which is in sharp contradiction with previous intuition, as *KAD* is assumed to provide a higher level of anonymity. The plot also shows that values obtained for the two *eDonkey* servers are similar, which indicates that very different filtering policies have no significant influence on the amount of paedophile queries.




Figure 1: Fraction of queries of each kind in our three datasets.

In order to gain a more detailed insight on this phenomenon, we study the frequencies of each keyword separately in the three datasets. As we want to explore possible correlations between the paedophile nature of a keyword and its frequency, we need a way to quantify the paedophile nature of a keyword. To do so, we use the 28-week dataset and the paedophile query detection tool presented in [6], which divides a dataset between *paedophile* and *not paedophile* queries. We denote by Q the whole dataset of queries, and by $Q(k)$ the set of queries containing a given keyword k . For each keyword

k , we obtain $Q(k) = N(k) + P(k)$, where $N(k)$ and $P(k)$ are the subset of queries containing keyword k and tagged as *not paedophile* or *paedophile*, respectively. We then define the *paedophile coefficient* $\pi(k)$ of keyword k as: $\pi(k) = \frac{|P(k)|}{|Q(k)|}$. If all the queries with keyword k are paedophile queries, $\pi(k) = 1$, and if none of them are, $\pi(k) = 0$. All keywords in the *not paedophile* category have a *paedophile coefficient* below 0.006. For keywords in the *mixed* category, the *paedophile coefficient* is above 0.01 and below 0.4. All *paedophile* keywords but one have a *paedophile coefficient* above 0.85. Finally, we plot in Figure 2 the ratios $\frac{f_{eDonkeyFR}(k)}{f_{kad}(k)}$ and $\frac{f_{eDonkeyUA}(k)}{f_{kad}(k)}$, where $f_s(k)$ denotes the frequency of queries composed of keyword k in the dataset s , for each of our 72 keywords. We rank keywords on the horizontal axis in increasing order of *paedophile coefficient*. The horizontal line represents $y = 1$, which enables a visual comparison of the values: if the point is below the line, then the keyword is more frequent in *KAD*, otherwise it is more frequent in the *eDonkey* dataset.

This plot gives a clear evidence for a correlation between the paedophile nature of a keyword and its higher presence in *eDonkey* than in *KAD*. In addition, the frequencies in both *eDonkey* datasets are very similar for the vast majority of keywords.

We therefore conclude that anonymity is not the prevailing factor when paedophile users choose a network, since neither the decentralised architecture of *KAD* nor the different filtering policies increase the frequency of paedophile queries. Instead, the frequency of paedophile queries is even higher in *eDonkey* than in *KAD*. Finding an explanation for this unexpected phenomenon is still an open question. The higher technical skills required to use *KAD* may be part of the explanation. Users may also search content on *eDonkey* while protecting their privacy with other tools, such as Virtual Private Networks or TOR [9]. The fact that in *KAD* search requests are sent over UDP and cannot benefit from TOR anonymisation could explain the difference in the network usage.

4 Ages indicators in queries

A way to gain more insight on observed paedophile activity is to study the distribution of age indicators in queries [8]. Notice that age indicators are sometimes used in other contexts than paedophile activity, especially when parents seek content suitable for children of a certain age. However, one can observe on Figure 2 that ages indicators have similar behavior to those obtained for the *paedophile* group, and are therefore closely related to the topic.




Figure 2: Ratio of keyword frequencies in *eDonkey* vs *KAD*. Keywords are ranked in increasing order of *paedophile coefficient*. Points above the $y = 1$ horizontal line indicates keywords more frequent in the corresponding *eDonkey* dataset; below the line keywords are more frequent in *KAD*.

We plot the distribution of age indicators on Figure 3: for each integer n lower than 17, we plot the number of queries of the form nyo in each dataset (yo stands for *years old*). The three plots have similar shape, with mostly increasing values from 1 to 10, a little drop at 11, a peak at 12 and a fall from 13 to 16. These values for *KAD* are below the values for the *eDonkey* servers, which is due to the fact that this dataset is a bit smaller than others and that paedophile queries are rarer in it. The key point here is that the distributions are very similar in all three datasets. This indicates that, although the *amount* of paedophile activity varies between systems, its nature is similar, at least regarding ages.

5 Quantifying paedophile activity in *KAD*

In [6], the authors establish a method to quantify the fraction of paedophile queries in *eDonkey*. It relies on a tool able to accurately tag queries as paedophile or not, and on an estimate of the error rate of this tool. Such an approach cannot directly be applied to *KAD* though, as only a small (and biased) fraction of all queries may be observed in this system. We however show in this section how to derive the fraction of paedophile queries in *KAD*




Figure 3: Distribution of age indicators in our three datasets.

from the one in *eDonkey*.

In a given system, *eDonkey* or *KAD* here, we consider different sets of queries and we denote by Q the set of all queries, P the subset of paedophile queries in Q , \overline{Q} the subset of queries composed of one word among the 72 monitored keywords, \overline{P} the subset of paedophile queries with one word, *i.e.* consisting of one of the 19 monitored paedophile keywords (and so: $\overline{P} = \overline{Q} \cap P$). Figure 4 illustrates our notations.




Figure 4: The different sets of queries we define for each considered dataset.

In both our *eDonkey* measurements, $|P|$ and $|Q|$ may be directly estimated [6] and one can then obtain the fraction $\frac{|P|}{|Q|}$ of paedophile queries in the dataset. We give the results for our two measurements in Table 1. On the contrary, in *KAD*, one may only estimate $|\overline{P}|$ and $|\overline{Q}|$.

However, we define $\alpha = \frac{|\overline{Q}| - |\overline{P}|}{|\overline{Q}| - |\overline{P}|}$ and $\beta = \frac{|\overline{P}|}{|\overline{P}|}$, which capture the probability for a non paedophile query, respectively paedophile, to make a query of one word among one of our monitored keywords. Given the definition of α and

dataset	$\frac{ P }{ Q }$	$ \overline{P} $	$ \overline{Q} $	α	β
<i>edonkeyFR</i>	$2.554 \cdot 10^{-3}$	74,557	241,152	$1.431 \cdot 10^{-3}$	0.2502
<i>edonkeyUA</i>	$2.668 \cdot 10^{-3}$	46,763	166,154	$1.538 \cdot 10^{-3}$	0.2251
<i>KAD</i>	n/a	30,821	250,000	n/a	n/a

Table 1: Main features of the three datasets.

β , there is no *a priori* reason to assume that they have significantly different values between *eDonkey* and *KAD*. From the definitions of α and β , we have:

$$\alpha = \frac{|\overline{Q}| - |\overline{P}|}{|Q| - |P|} \implies |Q| = |P| + \frac{|\overline{Q}| - |\overline{P}|}{\alpha} = \frac{\alpha|P| + |\overline{Q}| - |\overline{P}|}{\alpha}$$

$$\beta = \frac{|\overline{P}|}{|P|} \implies |P| = \frac{|\overline{P}|}{\beta}$$

Then, the following expression holds:

$$\begin{aligned} \frac{|P|}{|Q|} &= \frac{|\overline{P}|}{\beta} \times \frac{\alpha}{\alpha|P| + |\overline{Q}| - |\overline{P}|} \\ &= \frac{\alpha|\overline{P}|}{\beta|\overline{Q}| + (\alpha - \beta)|\overline{P}|} \end{aligned} \quad (1)$$

We now use expression (1) to infer the fraction of paedophile queries that were submitted in the *KAD* P2P network during our experiment. Using the values from Table 1 and the average values of α and β between our *eDonkey* datasets, we obtain:

$$\frac{|P|}{|Q|} \approx 0.087\% \pm 0.008$$

This value is of similar magnitude to the one of *eDonkey* (approx. 0.25%) but close to three times lower.

Notice that this estimation of $\frac{|P|}{|Q|}$ relies on the value of α . One may wonder whether the choice of keywords from which we built $\overline{Q} \setminus \overline{P}$ has a significant impact on the estimated value of $\frac{|P|}{|Q|}$ in *KAD*. We check this as follows: we randomly select 1,000 subsets of 26 keywords out of the 53 keywords which compose the queries in $\overline{Q} \setminus \overline{P}$. We then compute, for each subset, the number of queries consisting of exactly one of those keywords and the resulting value of alpha. For *edonkeyFR*, we obtain an average value of $\bar{\alpha} = 0.000889$ (minimum: 0.000256, maximum: 0.00153, and 90% of the values in [0.000463;0.00133]). For *edonkeyUA*, we obtain an average value

of $\bar{\alpha} = 0.00105$ (minimum: 0.000352, maximum: 0.00172, and 90% of the values in [0.00062;0.00148]). This means that we would obtain very similar results with 26 keywords only and so we may be confident in our estimate obtained with 53 keywords.

6 Conclusion

We performed a first comparative study of two large-scale peer-to-peer networks, *KAD* and *eDonkey*, with regards to the queries related to child pornography. We designed a methodology to collect and process datasets allowing to compare them in a relevant manner. We obtained the counter-intuitive result that paedophile keywords are significantly more present in *eDonkey* than in *KAD*, despite the higher anonymity level it provides. On the contrary, our study of age indicators in queries showed that the nature of paedophile queries is similar in these systems. We finally established the first estimate of the fraction of paedophile queries in *KAD*. We obtained a value close to 0.09%, which is of the same magnitude but significantly lower than in *eDonkey* (0.25%).

Our contributions open various directions for future work. In particular, our methodology may be applied to compare other systems, and our datasets may be used to perform either deeper analyses on paedophile activity or on general search engine behaviors.

Acknowledgments

This work is supported in part by the MAPAP SIP-2006-PP-221003 and ANR MAPE projects.

References

- [1] Frédéric Aidouni, Matthieu Latapy, and Clémence Magnien. Ten weeks in the life of an eDonkey server. *International Workshop on Hot Topics in P2P Systems*, 2009.
- [2] Thibault Cholez, Isabelle Chrisment, and Olivier Festor. Monitoring and Controlling Content Access in KAD. In *International Conference on Communications - ICC 2010*, Capetown, Republic of South Africa, May 2010. IEEE.

- [3] European Commission. Safer Internet programme: Empowering and protecting children online. http://ec.europa.eu/information_society/activities/sip/index_en.htm.
- [4] Oliver Heckmann, Axel Bock, Andreas Mauthe, and Ralf Steinmetz. The eDonkey file-sharing network. In Peter Dadam and Manfred Reichert, editors, *GI Jahrestagung (2)*, volume 51 of *LNI*, pages 224–228. GI, 2004.
- [5] Daniel Hughes, James Walkerdine, Geoff Coulson, and Stephen Gibson. Is deviant behavior the norm on P2P file-sharing networks? *IEEE Distributed Systems Online*, 7(2), 2006.
- [6] Matthieu Latapy, Clémence Magnien, and Raphaël Fournier. Quantifying paedophile activity in a large P2P system. *Information Processing & Management*, 2012.
- [7] Declan McCullagh. RIAA: Child porn rife on P2P networks. *CNet*, 2003. http://news.cnet.com/RIAA-Child-porn-rife-on-P2P-networks/2100-1028_3-5073817.html.
- [8] Chad M.S. Steel. Child pornography in peer-to-peer networks. *Child Abuse & Neglect*, 2009.
- [9] The Tor Project, Inc. TOR project. <https://www.torproject.org/>.