Removing bias due to finite measurement of dynamic systems: case study on P2P systems

Lamia Benamara[1] and Clémence Magnien[1]

[1]*LIP6, CNRS & Université Pierre et Marie Curie - 4 place Jussieu, 75005 Paris, France - Prenom.nom@lip6.fr*

Mesurer avec précision la dynamique des graphes de terrain est une tâche difficile, car les propriétés observées peuvent être biaisées pour différentes raisons, en particulier le fait que la période de mesure soit finie. Dans ce papier, nous introduisons une méthodologie générale qui nous permet de savoir si la fenêtre d'observation est suffisamment longue pour caractériser une propriété donnée dans n'importe quel système dynamique.

Nous appliquons cette méthodologie à l'étude des durées de sessions et des durées de vie des fichiers sur deux jeux de données P2P. Nous montrons que le comportement des propriétés est différent : pour les durées de sessions, notre méthodologie nous permet de caractériser avec précision la forme de leur distribution. Par contre, pour les durées de vie des fichiers, nous montrons que cette propriété ne peut pas être caractérisée, soit parce qu'elle n'est pas stationnaire, soit parce que la durée de notre mesure est trop courte.

1 Introduction

Many systems are naturally dynamic. For instance in the internet, routers, AS and/or links between them are created or deleted [**?**]; in peer-to-peer (P2P) networks users join or leave the system [**?**, **?**, **?**] and exchange different files at different times. In all these cases, understanding the dynamics of the system is a key issue. However, accurately measuring these dynamics is a difficult task. In particular, the fact that the observation window is necessarily finite induces a bias for property characterization [**?**, **?**]. Though this bias tends to decrease when the observation window length increases, it is difficult to quantify it in practice, or know whether it is negligible or not.

In this paper, we introduce a new methodology that allows to rigorously determine the minimum observation time required to characterize a stationary property in real-world dynamic systems. This methodology is different and complementary to other methodologies existing in the literature [**?**, **?**, **?**], and has two main advantages. First, it allows to determine if the observation window was long enough for a rigorous characterization. Second, it can be applied to any property characterizing the dynamics of a system. To illustrate its relevance, we apply it to the study of session lengths and files' life duration in two different P2P systems.

2 Methodology

Suppose we start observing a dynamic graph at a time t, for a duration l. We denote by $W_{t,l}$ this observation window. We are faced with two problems if we want to characterize the graph's dynamics from the observation of $W_{t,l}$. First, l must be long enough for $W_{t,l}$ to be *representative*. Second, even if it is representative, the fact that l is *finite* still induces a bias for property characterization. Indeed, events occurring before t or after $t+l$ are not observed, which prevents from characterizing accurately some quantities. An important point to observe is that the longer the measurement period, the smaller the bias induced.

* We thank Pierre Borgnat, Patrice Abry, Matthieu Latapy, Marcelo Dias de Amorim and Nadjet Belblida for valuable ideas and comments about this manuscript. This work was supported in part by the French ANR MAPE project, and by a grant from the *Agence Nationale de la Recherche*, with reference ANR-10-JCJC-0202. This research work is conducted and funded by the European Commission through the EULER project (Grant No 258307) part of the Future Internet Research and Experimentation (FIRE) objective of the Seventh Framework Programme (FP7).

Our methodology addresses these two issues at the same time. Intuitively, it aims at deciding if the measurement period $W_{t,l}$ is long enough to characterize a given property P, i.e. if the bias induced by its finiteness on the observed property is negligible. If the window $W_{t,l}$ is long enough, then if we use a longer window $W_{t,l+x}$, the observed property does not change: $P(W_{t,l}) = P(W_{t,l+x})$. In order to know if a given window is long enough, we use windows of increasing length W_{0,l_1},W_{0,l_2}, ... W_{0,l_n}, with $l_1 < l_2 < ... < l_n$. By studying how the observed property $P(W_{0,l_1})$,$P(W_{0,l_2})$, ...$P(W_{0,l_n})$ evolves as a function of l, we determine if it is correctly evaluated or not.

Finally, an important point is that characterizing a property P only makes sense if it is stationary, i.e. if P does not evolve while the measurement is under progress. Notice however that if it is not stationary, our methodology will not be able to provide a characterization: the observed property P will not become stable when the observation window length l increases. If it does become stable, this means both that $W_{t,l}$ is long enough, and that P is stationary. Notice that, depending on the property studied, other types of bias can occur, see for instance [**?**], including biases coming from the identification of users and their sessions. We will also rigorously take this into account, see Section **??**.

Here, most of the properties we study are complementary cumulative distributions, i.e. for each value k, P_k is the fraction of all observations values which are larger than or equal to k.

To study how an observed distribution P evolves with the length of the observation window, we will first plot the observed distributions $P(W_{t,l})$ for different values of l. In order to confirm more formally the visual observations, we will also study a statistical indicator which quantifies how close two distributions P and Q are to each other: the *Monge-Kantorovich distance*, or M-K distance [**?**] compares two normalized cumulative (complementary or not) distributions P and Q. It is equal to the mean of the distance between the two distributions: $MK(P,Q) = (\sum_k |P_k - Q_k|)/k_{\max}$.

We use this indicator to study how the observed distribution $P(W_{t,l})$ evolves: we compute the M-K distance between $P(W_{0,l})$ (with different values of l) and $P(W_{0,l_{max}})$, where l_{max} is the length of the longest observation window for this dataset, and plot this as a function of l. Following [**?**], we also study the mean and the standard deviation of $P(W_{0,l})$ as a function of l.

3 Data

In order to show the relevance of our methodology, we use two datasets: the *queries* dataset which is a capture of the UDP traffic of a large *eDonkey* server [**?**]. It consists of the queries made by users (for lists of files matching certain keywords, or for providers for a given file), and of the server's answers to these queries. The measurement lasted for 10 weeks which represents 1 billion messages, with 89 million peers and 275 million files involved. The *logins* dataset consists in a trace of the login and logout of peers on the *eDonkey* network [**?**]. It contains more than 200 millions of connections by more than 14 millions of peers, over a period of 27 days. The two datasets are therefore complementary.

4 Users' session lengths

4.1 Definition of a session

We do not formally know when user sessions begin or end in the *queries* dataset, because there is no notion of session in the UDP *eDonkey* protocol. Instead, users make stand-alone queries and receive answers from the server. We therefore have to infer sessions from these queries.

It is natural to consider that two consecutive queries made by a same user belong to the same session (whether they are for a same file or not) if the time elapsed between them is short, and belong to two different sessions if it is long. The question is then to find an appropriate threshold for distinguishing between these two cases. Based on the study of the inter-query time distribution (not presented here), we have chosen to use a threshold of 10 800 seconds, i.e. 3 hours.

4.2 Characterization of session lengths

We now apply our methodology to the study of the session length distributions S, by studying $S(W_{0,l})$ for different values of l.

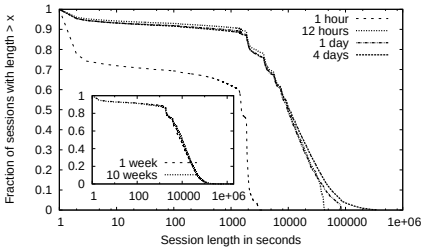

Figure 1: Complementary cumulative distributions of $S(W_{0,l})$ for different observation windows lengths in log-lin scale, for the *queries* dataset.

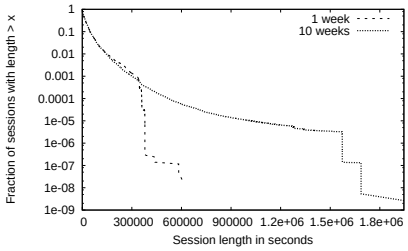

Figure 2: Complementary cumulative distributions of $S(W_{0,l})$ for observation windows lengths $l = 1$ week and $l = 10$ weeks in lin-log scale, for the *queries* dataset.

Figure **??** shows the complementary cumulative distribution $S(W_{0,l})$ for different values of l, up to $l =$ 10 weeks, for the *queries* dataset. The shapes of these distributions are similar, with a small fraction of sessions with length smaller than 2 000 s, and an approximately linear shape between 2 000 s and 100 000 s. However, when $l \leq 1$ day, the distributions exhibit a clear cut-off. This is not the case anymore for $l \geq 4$ days: the tail of the distribution flattens after a bend occurring close to 100 000 s ($\sim$ 28 hours), and we observe a small fraction of *extreme* values after this bend. For observation windows larger than four days, the shape of the distribution does not seem to evolve anymore: the distributions corresponding to $l = 1$ week and $l = 10$ weeks (presented in the inset) are very similar to each other and to the one obtained for $l = 4$ days.

One must be however careful when driving conclusions from a visual examination. Indeed, if we observe the same plot as the inset of Figure **??** but with a linear scale on the *x*-axis and a logarithmic scale on the *y*-axis (see Figure **??**), the distributions seem visually strongly different from each other. However, the distributions are different only for less than 1% of the values, which are values after the bend in Figure **??** and are *extreme* values. The fact that the extreme values change when l increase shows that they cannot be characterized with our methodology, and we leave their study for further work.

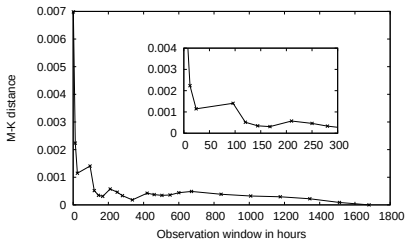

Figure 3: MK($S(W_{0,l})$, $S(W_{0,l_{max}})$) as a function of l, for the *queries* dataset.

To confirm these observations, we study MK($S(W_{0,l})$, $S(W_{0,l_{max}})$) as a function of l, presented in Figure **??**. The values observed tend to decrease (with fluctuations) until the observation window reaches approximately 150 hours (6 days and 6 hours). After this, the value of the M-K distance becomes very small: this shows that the corresponding distributions are very close to each other.

We also studied the standard deviation and the mean of $S(W_{0,l})$ as a function of l (not presented here). We observe that the mean becomes stable once l reaches approximately 1 week, at the same time as the M-K distance. This confirms that an observation window of one week is long enough to accurately estimate the distribution. The standard deviation, however, does not seem to converge as the observation window length increases, confirming that the distribution cannot be *fully* characterized. This is consistent with the distinction between the normal part of the distribution and extreme values.

Figure **??** shows the complementary cumulative distribution $S(W_{0,l})$ for different values of l, up to $l = 3$ weeks, for the *logins* dataset. We can see that the shape of these distributions are similar, and get closer to each other as l increases. However, when we compare these distributions with the M-K distance (see Figure **??**), the values obtained tend to decrease linearly which means that the distributions change at a constant rate. The values obtained for the mean and the standard deviation also do not stabilize. Therefore, we can not fully characterize this distribution. We however have confidence that the true shape of the distribution is not far from the one we observed.

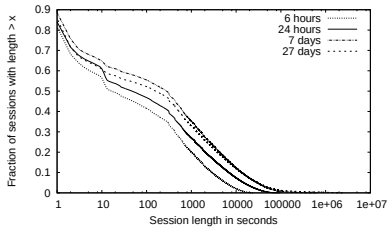

Figure 4: Complementary cumulative distributions of $S(W_{0,l})$ for different observation windows lengths, for the *logins* dataset.

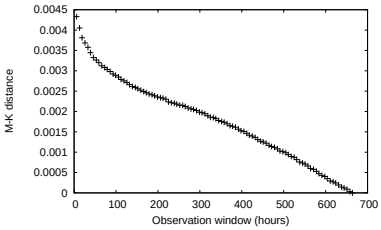

Figure 5: MK($S(W_{0,l})$, $S(W_{0,l_{max}})$) as a function of l, for the *logins* dataset

5 Files' lifetime

We considered two different definitions for a files' lifetime F. The first one is the same as for users' sessions lengths: we use a threshold and consider that a file is not present in the system if there is no consecutive queries for this file separated by less than this threshold. The second definition consists in considering the time interval between the first and the last query for a given file. In both cases, the shape of distributions $F(W_{0,l})$ (not presented here) evolves strongly with l. We therefore conclude that this property cannot be characterized. The question which arises is whether this is because this property is intrinsically not stationary or because our measurement period is too short to be able to characterize it.

6 Conclusion

In this paper we introduced an empirical methodology for deciding when the bias induced by the finiteness of observation windows in dynamic systems becomes negligible. To illustrate the relevance of this approach, we applied it to the study of sessions lengths and files' life duration in two different datasets.

We have shown that we can characterize some properties, but not all. When a property can't be characterized, our methodology doesn't allow to determine if the observation window shall be increased or not since we don't know the stationarity of the property itself. It is interesting to note that, for a same dataset, some properties can be accurately characterized, and others not.

References

[ALM09] Frédéric Aidouni, Matthieu Latapy, and Clémence Magnien. Ten weeks in the life of an edonkey server. In *Proceedings of HotP2P'09*, 2009.

[GKT09] Tryphon T. Georgiou, Johan Karlsson, and Mir Shahrouz Takyar. Metrics for power spectra: An axiomatic approach. *IEEE Transactions on Signal Processing*, 57(3):859–867, 2009.

[GT99] Matthias Grossglauser and David N. C. Tse. A framework for robust measurement-based admission control. *IEEE/ACM Trans. Netw.*, 7(3):293–309, 1999.

[LBFM09] Stevens Le-Blond, Fabrice Le Fessant, and Erwan Le Merrer. Finding good partners in availability-aware P2P networks. In *SSS*, pages 472–484, 2009.

[MOVL09] Clémence Magnien, Frédéric Ouédraogo, Guillaume Valadon, and Matthieu Latapy. Fast dynamics in internet topology: Observations and first explanations. In *2009 Fourth International Conference on Internet Monitoring and Protection*, pages 137–142. IEEE, 2009.

[SGG03] Stefan Saroiu, Krishna P. Gummadi, and Steven D. Gribble. Measuring and analyzing the characteristics of napster and gnutella hosts. *Multimedia Systems*, 9:170–184, 2003.

[SR06] Daniel Stutzbach and Reza Rejaie. Understanding churn in peer-to-peer networks. In *Internet Measurement Conference*, pages 189–202, 2006.

[WAL04] Walter Willinger, David Alderson, and Lun Li. A pragmatic approach to dealing with high-variability in network measurements. In *Internet Measurement Conference*, 2004.