Estimating properties in dynamic systems: the case of churn in P2P networks

Lamia Benamara
UPMC Univ Paris 06, UMR 7606,
LIP6, F-75016, Paris, France
CNRS, UMR 7606, LIP6, F-75016, Paris, France
Email: lamia.benamara@lip6.fr

Clémence Magnien
UPMC Univ Paris 06, UMR 7606,
LIP6, F-75016, Paris, France
CNRS, UMR 7606, LIP6, F-75016, Paris, France
Email: clemence.magnien@lip6.fr

***Abstract*—**
In many systems, such as P2P systems, the dynamicity of participating elements, or *churn*, has a strong impact. As a consequence, many efforts have been made to characterize it, and in particular to capture the session length distribution. However in most cases, estimating it rigorously is difficult. One of the reasons is that, because the observation window is by definition finite, parts of the sessions that begin before the window and/or end after it are missed. This induces a bias. Although it tends to decrease when the observation window length increases, it is difficult to quantify its importance, or how fast it decreases.

Here, we introduce a general methodology that allows us to know if the observation window is long enough to characterize a given property. This methodology is not specific to one study case and may be applied to any property in a dynamic system. We apply this methodology to the study of session lengths in a massive measurement of P2P activity in the *eDonkey* system. We show that the measurement needs to last for at least one week in order to obtain representative results. We also show that our methodology allows us to precisely characterize the shape of the session length distribution.

I. Introduction.

Many systems are naturally dynamic. For instance in the internet, routers, AS and/or links between them are created or deleted [10], [11]; in P2P networks users join or leave the system [15], [13], and exchange different files at different times; in online social networks users may create or delete accounts, or cease to be active, and create or delete connections with other users [16].

In all these cases, understanding the dynamics of the system is a key issue. However, accurately measuring this dynamics is a difficult task. In particular, the fact that the observation window is necessarily finite induces a bias in the observations [12], [15], [13]. Though this bias tends to decrease when the observation window length increases, it is difficult to quantify it in practice, and know whether it is negligible or not.

Another problem is that a small observation window may not be representative of the whole behavior of the system. For instance, measuring the activity in a P2P system during one hour is not enough to capture fully the dynamics of user usages, because of day/night activity variations. However, it is not *a priori* clear whether one day, or two, or one week, is long enough.

In this paper, we introduce a new methodology that allows to rigorously characterize dynamic metrics in real-world dynamic systems. This methodology is different and complementary to other methodologies existing in the literature [15], [13], and has two main advantages:

- it allows to determine if the observation window length was sufficient for a rigorous characterization;
- it can be applied to any property characterizing the dynamics of a system.

To illustrate the relevance of this methodology, we apply it to the study of session lengths in a large P2P system, which captures the dynamicity of users, or *churn*.

This document is organized as follows. In Section II, we introduce our methodology and present our dataset, as well as the choices we made to identify sessions in this dataset. In Section III we present the application of our methodology to the study of session lengths. We present related work in Section IV, and our conclusions and future work in Section V.

II. Methodology and Data

A. Methodology

When trying to characterize the dynamics of a system, one is faced with two problems. First, the observation window must be long enough to be *representative*. Second, even if it is representative, the fact that it is *finite* still induces a bias in the observations.

Events occurring before or after the observation window are not observed, which prevents from characterizing accurately some quantities (for instance, session lengths, or time correlations between different events).

Our methodology addresses these two issues at the same time. Intuitively, it aims at using an observation window long enough, so that the bias induced by its finiteness on the observed property becomes negligible. However, it is not possible to know *a priori* how long this should be.

If the window is long enough, then if we use a longer window, the value of the observed property will be the same. However, once the measurement is completed, it is not possible to increase the observation window length. We therefore extract windows that are *shorter* than the actual observation window. We then compare the properties obtained for different lengths of the observation window. By studying how properties evolve as a function of the window length, we determine if they are correctly evaluated or not: if a property fluctuates or varies greatly as the window length increases, then the property is certainly not accurately evaluated. Indeed, a slightly shorter or longer observation window would have yielded a different value. Instead, if an observed property tends to become stable as the window length increases, then it is probably accurate.

Finally, an important point is that characterizing a property only makes sense if the system is stationary, i.e. if this property does not evolve while the measurement is under progress. Notice however that if the system is not stationary, our methodology will not be able to provide a characterization: the observed property will not become stable when the observation window length increases. If it does become stable, this means both that the observation window is long enough, and that the property is stationary [1]. Here, we study distributions, therefore our methodology is applied by studying different distributions for different sizes of the observation window. After this, we use standard statistical tools to compare them.

B. Data

We applied this methodology to the study of session lengths in a P2P system. The data we use comes from [1], and consists in the 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, of for providers for a given file), and of the server's answers to these queries. The measurement lasted for 10 weeks. This represents 1 billion messages, with 89 million peers and 275 million files involved. This dataset is publicly available [1].

We are interested in the session lengths of users. Different types of bias can occur in this case. Stutzbach and Rejaie [15] gave a comprehensive list of the most common of these biases. Because our data consists of *all* traffic managed by the server during the measurement period, many of these biases do not occur in our case: we observe all users and all messages.

The first bias that occurs in our context, which we already discussed, comes from the fact that the observation window is finite, which is central to this paper. Other biases comes from the identification of users and their sessions. We detail them below.

First, identifying users is a difficult question. We only have access to the IP addresses of the computers from which queries are entered. Computers are identified by an IP address at a given time, but this may change and we are unable in general to detect that a same computer has two different addresses (because of dynamic addresses for instance) and/or that two computers are using the same address (because they are behind a same NAT for instance). In addition, a same user may use several computers, and several users may use the same computer, making identification of users even more challenging. In the absence of a satisfying method for identifying users, we decided here to assume that each IP address corresponds to one user. Improving this is one of our main perspectives.

[1]Note that the system may be stationary with respect to some properties and not others; in such a case our methodology may provide a characterization for the stationary properties and not the others.

C. Definition of a session

In order to apply our methodology to the study of session lengths, we have to identify sessions in our data. However, we do not formally know when sessions begin or end, because there is no notion of session in the UDP *eDonkey* protocol. Instead, users make stand-alone queries and receive answers from the server.

Therefore, we infer sessions for a given IP address by studying the time elapsed between consecutive queries. It is natural to consider that two consecutive queries made by a same IP address belong to the same session 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.

To study this, we computed the inter-query time distribution, presented in Figure 1 (we display both the distribution (a) and the complementary cumulative distribution (b)). We observe clear peaks at 60 second and at multiples of it (120, 240, 300, 900 s, ...) in the distribution (they can be more clearly seen in the inset). These peaks indicate that, though users decide which queries to make and when they make them, there is a strong influence of the protocol on the observed data: most client applications automatically perform periodical queries. Although these peaks become smaller after 1800 s, a zoom on the plot (not presented here) shows that they are clearly defined for values up till at least 20 000 seconds.

To study the importance of these peaks, we computed, for a same measurement window, the session length distributions obtained with two different thresholds, the first chosen just before a peak and the second just after this same peak. We made a comparison between these two distributions and we observed no significant difference.

In order to smooth out the plot, we consider the complementary cumulative distribution (Figure 1 (b)). There is a high density of values between 1 000 and a value slightly smaller than 10 000 (the slope of the distribution is steep in this region). Such a high density indicates normal inter-query lengths within a session, and choosing a threshold in this region or before it would have little meaning. Therefore, we argue that the threshold must be at least as large as 10 000 seconds.

We have chosen, for all the rest of the paper, to use a threshold of $t = 10\,800$ seconds, i.e. 3 hours. Therefore, in all the paper, if a same IP address sends consecutive queries separated by less than three hours, these queries belong to a same session, otherwise they belong to different sessions [2].

[2]A detailed study of session lengths would probably benefit from studying other values for this threshold. However our goal in this paper is to illustrate our methodology and show that we can obtain interesting insight on the characteristics of session lengths. Other thresholds lead to similar results.

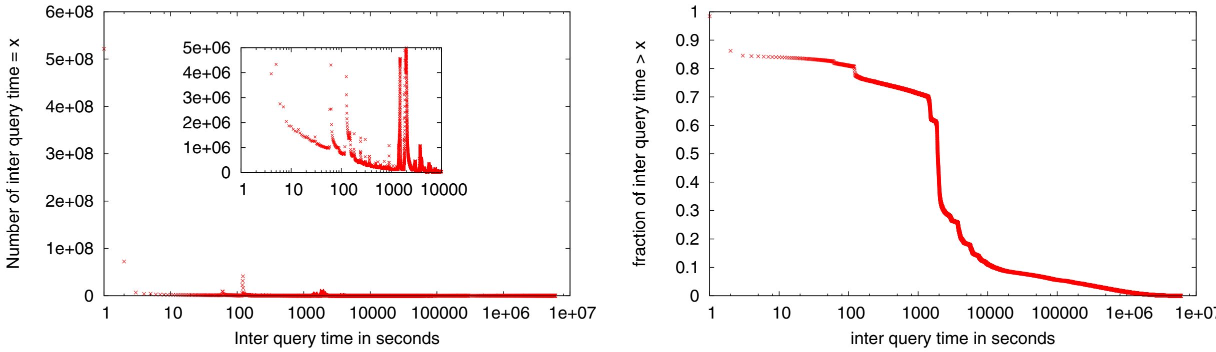

(a) Distribution.

(b) Complementary cumulative distribution.

Figure 1: Inter-query time distribution.

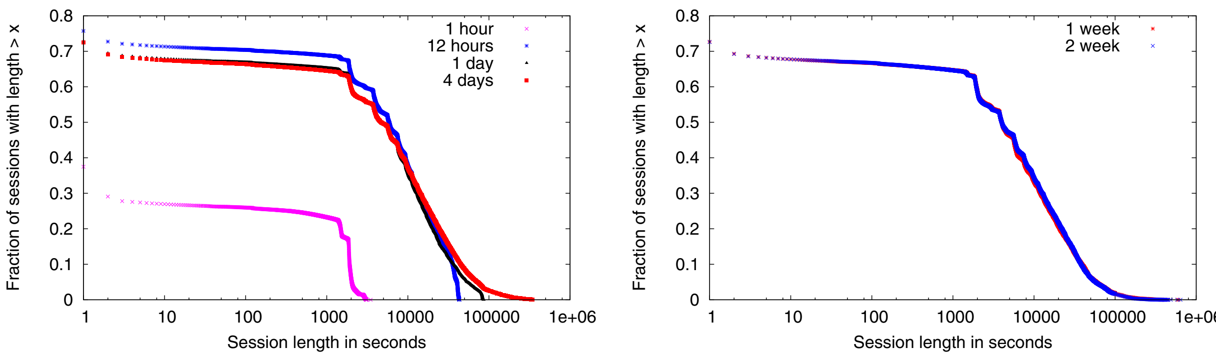

(a) 1h, 12h, 1 day and 4 days.

(b) 1 and 2 weeks (the two distributions overlap almost completely).

Figure 2: Complementary cumulative distributions of session lengths for different observation windows lengths.

III. Characterization of Session Lengths

We now apply our methodology to the study of the session length distribution. First, these distributions are highly irregular. They present clear peaks and valleys, which are linked to the peaks in the inter-session time distribution, see Figure 1 (a). Similar observations hold for different observation window lengths and positions. In this section, we will therefore consider complementary cumulative distributions, to smooth out the irregularities.

A. *Study of distribution*

Figure 2 shows the complementary cumulative distribution of session lengths for different observation window lengths. The fractions of sessions with length 0 are not the same, which causes the normalized distributions to be vertically shifted[3]. The shapes of these distributions are however 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, for measurement windows shorter than or equal to one day, the distributions exhibit a clear cut-off. This is not the case anymore after four 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.

[3]Since the x-axis is in log-scale, the dot $(0, 1)$ which belongs to all these distributions does not appear.

For observation windows larger than four days, the shape of the distribution does not seem to evolve anymore: Figure 2 (b) shows that the distributions for one and two weeks are very similar to each other and to the one obtained for four days.

However, when the length of the observation window increases, we again observe a small difference between the corresponding distributions. Figure 3 shows the distributions obtained for windows of 1 and 10 weeks. We observe in Figure 3 (a) a small gap between them, caused by the fraction of sessions of length 0 (which does not appear because of the log-scale on the x-axis): when the distribution is normalized by the number of sessions with length strictly larger than 0 (Figure 3, b), this gap disappears. This shows that, though the shape of the distribution does not vary anymore, the fraction of sessions with length 0 does.

We observed the same phenomena with distributions obtained for observation windows of the same length, but located at different positions in the whole measurement window.

This shows that the fraction of sessions 0 changes with the measurement period but that the general form of the distribution, when these values are not taken into account, does not change.

We saw that the distributions seem *visually* not to change

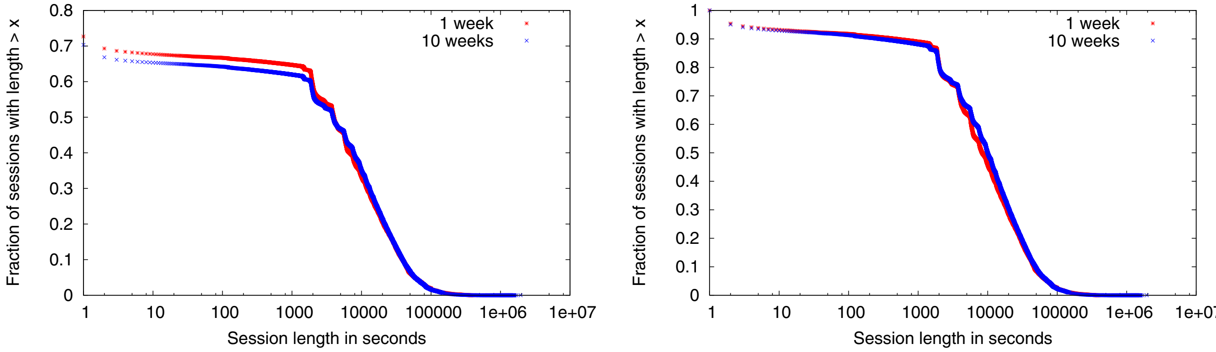

(a) Normalized distribution.

(b) Distribution normalized by number of values larger than 0.

Figure 3: Complementary cumulative distributions of session lengths for observation windows of 1 and 10 weeks.

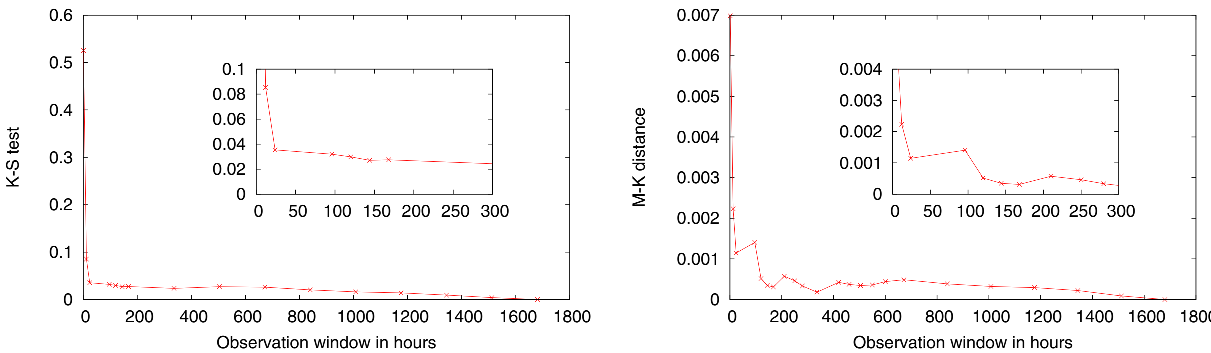

(a) K-S test.

(b) M-K distance.

Figure 4: Comparison of the distributions obtained for observation windows of 1,2,3,....10 weeks to the distribution obtained for the whole 10 weeks of the measurement window, using the K-S test and the M-K distance.

once the observation window length has reached four days. However, one must be careful when driving conclusions from a visual examination. Indeed, Figure 5 shows the same distributions, but with a linear scale on the x-axis and a logarithmic scale on the y-axis. At first glance, the distributions seem strongly different from each other. However, a more careful examination shows that the distributions are similar for at least 99% of the values. They are different only for values larger than approximately $150\,000$ s, which are values seen after the bend and are significantly rarer than values below this bend. This leads us to consider them as *extreme* values. This shows that the *normal* part of the distribution does not evolve anymore, and therefore is accurately characterized. The extreme values cannot be characterized in this way, and we leave their study for further work.

B. Statistical tests

In order to confirm more formally these observations, we study two indicators for how close two distributions are to each other.

The *Kolmogorov-Smirnov test*, or K-S test [4] compares two normalized cumulative (complementary or not) distributions. It is equal to the maximum, for all values k, of the distance between the two distributions: $\max_k |p(k)-p'(k)|$. It is always lower than 1, and the closer it is to 0, the more similar the two distributions are.

Figure 4 (a) presents the comparison with the K-S test of the session length distributions obtained for different measurement durations. A point with coordinate d on the x-axis represents the K-S test between the distribution for an observation window of length d and the distribution for the whole measurement period. The first values are high, and decrease quickly to approximately 4% for a measurement duration of 24 hours. After this, the decrease is linear. This clearly shows that measurement durations of less than 24 hours are not representative. However, we do not know if the value 4% is small enough to consider that the distributions are similar or not. Moreover, the linear shape does not correspond to a value which fluctuates before becoming stable. This plot does therefore not allow us to know when the observation window becomes long enough, or even to know if this happens during the measurement. Therefore, we cannot reach a conclusion with the K-S test.

The *Monge-Kantorovich distance*, or M-K distance [5], helps to answer these questions. It is equal to the mean of the distance between the two distributions: $(\sum_k |p(k) - p'(k)|)/k_{\max}$. Two distributions that only differ in a single point will therefore have a high K-S test, but a small M-K

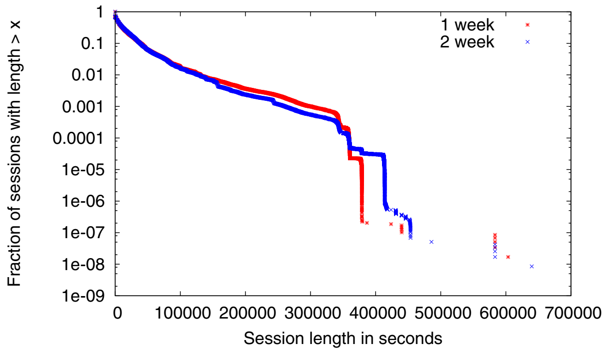

Figure 5: Complementary cumulative distributions of session lengths for observation windows of 1 and 2 weeks in lin-log scale.

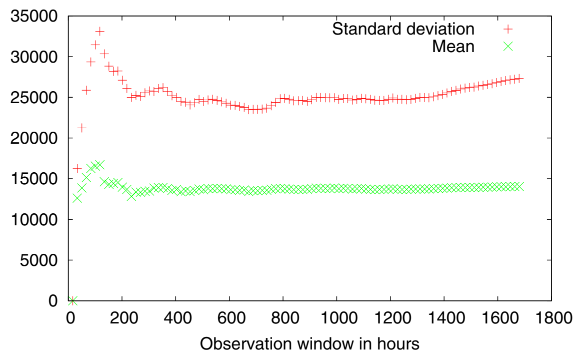

Figure 6: Mean and standard deviation of the session length distribution, as a function of the observation windows size.

distance.

Figure 4 (b) presents the difference between the same distributions, evaluated with the M-K distance. We observe a different behavior than with the K-S test: the values observed tend to decrease (with fluctuations), until the measurement period reaches 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.

Following [18], we also studied the standard deviation of the session length distribution (as well as its mean) as a function of the measurement window length. They are presented in Figure 6. We can see that the mean becomes stable after some time (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 [4], confirming that the distribution cannot be *fully* characterized. This is consistent with the opposition between the normal part of the distribution and extreme values. Indeed, the extreme values are very large and therefore have a strong impact on the standard deviation. The fact that they cannot be characterized causes the standard deviation to vary, whereas the fact that the normal part of the distribution is characterized causes the mean to become stable [5].

This confirms the intuition obtained by a visual study of the distributions: once the observation window length reaches one week, the normal part of the session length distribution stops evolving. This means two things. First, this distribution is stationary over time scales of the order of the whole measurement length, and it therefore makes sense to characterize it. Second, an observation window of one week is long enough to accurately estimate it. The extreme values of this distribution cannot however be characterized by our methodology.

IV. Related Work

Several papers have studied churn in P2P networks, see for instance [15], [14], [9], [3] and references within. These works have studied different properties related to churn, in particular session lengths, inter-session time, and correlations between consecutive session lengths. Some authors have also studied churn in other systems [6], [16].

Some authors have studied explicitly the bias induced by the measurement procedure on the observed churn. The authors of [2] acknowledge that changing the observation window length impacts the observed properties of churn. They do not however propose any specific methodology for dealing with this bias.

The *create-based method* [12], [13] removes the bias towards short sessions, created by the fact that it is only possible to observe the length of sessions that begin and end within the measurement window. To remove this bias, they propose to divide the measurement window of length T into two halves, and only consider sessions that begin during the first half. This leads to an unbiased estimation of sessions of length less than $T/2$. This methodology is complementary to the one we introduce here, which does not formally remove the bias, but allows to make observations for the shape of the distribution even for values larger than $T/2$.

Stutzbach and Rejaie [15] studied different aspects of peer dynamics in three different classes of P2P systems (Gnutella, Kad and BitTorrent). They carefully analyzed the different kinds of bias that may influence such a study, and presented a list of the bias they identified. They use the create-based method to deal with the bias caused by the finiteness of the observation window.

Wang *et al.* [17] argue that the create-based method is biased when the data is obtained through periodic sampling (which is not our case). They propose a new sampling algorithm called RIDE (ResIDual-based Estimator) which measures session length distributions with high accuracy and requires a low sampling frequency.

Finally, Willinger et al [18] also addressed, in the context of IP flows, the question of whether the observation window

[4]Notice that, if we had stopped the measurement at 1200 hours, we would have the impression that it converges, hence the importance to have an observation window large enough.

[5]When we study the same plot without taking into account the extreme values (i.e. cutting the distribution above a given value) the standard deviation tends to became stable.

is long enough to characterize some dynamic properties. They study the standard deviation of the flow size distribution as a function of the measurement length, and argue that the fact that it does not converge means that the samples may come from an underlying distribution with infinite variance. This in turn may make it difficult to fit the observed properties with a model. In our case, we showed that, even though the standard deviation does not converge, this is caused by extreme values, and we still are able to characterize the part of the distribution that does not contain these extreme values.

V. Conclusion and Future Work

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 in a large P2P system.

This showed that, in this case: (1) the observation window needs to be at least one week long in order to characterize session lengths; (2) it is possible to know the shape of the session length distribution independently of the observation window length (if it is longer than one week); this distribution is characterized by *normal* session lengths, and *extreme* values. While the distribution of the normal session lengths is independent of the observation window length, the extreme values elude statistical characterization. They are therefore dependent on the observation window length, but do not alter the *shape* of the distribution.

In this paper we applied our methodology to the study of churn in P2P systems. This methodology is however general, and can be applied to *any* property in a dynamic system. We believe that using it would yield interesting insights in different contexts, such as churn in other systems, or more generic studies of other dynamic systems, such as the internet, online social networks, etc.

Our observations about session lengths could be extended further. First, we aim at comparing our analysis with other *eDonkey* traces, in order to assess the generality of our observations. Second, we exhibited a distinction between normal and extreme session lengths, but we did not propose a method for formally distinguishing between them. In the same way, studying session lengths from a user point of view may lead to the discovery of different user classes, having different types of behavior in the system.

Studying models of user activity would allow us to gain a better intuition on the studied phenomena, and confirm formally our results. It may also provide *formal* bounds for the minimum observation window length needed to characterize the session length distribution with a given accuracy. Such results could then be extended to other dynamic properties.

Finally, we presented here a methodology for dealing with the bias introduced when measuring the dynamics of a system. In many systems, and in particular in the case of the internet, it is known that the measurement procedure may introduce a *structural* bias even if the system does not evolve with time. Some methods have been introduced to remedy this, see for instance [8], [7]. We believe it is therefore crucial to combine methodologies such as the one we introduced here, which deal with the dynamic bias, to methodologies dealing with the structural bias, in order to capture the properties of systems such as the internet, as well as their dynamics.

Acknowledgments

We thank Pierre Borgnat, 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 European MAPAP SIP-2006-PP-221003 project, the French ANR MAPE project, and the DyGram project funded by RNSC.

References

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

[2] Ranjita Bhagwan, Stefan Savage, and Geoffrey M. Voelker. Understanding availability. In *IPTPS*, pages 256–267, 2003.

[3] R. Bolla, R. Gaeta, A. Magnetto, M. Sciuto, and M. Sereno. A measurement study supporting p2p file-sharing community models. *Computer Networks*, 53:485–500, 2009.

[4] I. M. Chakravarti, R. G. Laha, and J. Roy. *Handbook of Methods of Applied Statistics*, volume I. John Wiley and Sons, USE, 1967.

[5] 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.

[6] Lei Guo, Enhua Tan, Songqing Chen, Xiaodong Zhang, and Yihong (Eric) Zhao. Analyzing patterns of user content generation in online social networks. In *ACM SIGKDD: International Conference on Knowledge Discovery and Data Mining*, pages 369–378, 2009.

[7] A. Lakhina, J. Byers, M. Crovella, and P. Xie. Sampling biases in IP topology measurements. In *IEEE INFOCOM*, 2003.

[8] Matthieu Latapy and Clémence Magnien. Complex network measurements: Estimating the relevance of observed properties. In *2008 IEEE INFOCOM - The 27th Conference on Computer Communications*, pages 1660–1668. IEEE, 2008.

[9] Thomas Locher, David Mysicka, Stefan Schmid, and Roger Wattenhofer. A peer activity study in edonkey and kad. In *International Workshop on Dynamic Networks: Algorithms and Security (DYNAS)*, 2009.

[10] 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.

[11] R. Oliveira, B. Zhang, and L. Zhang. Observing the evolution of internet AS topology. In *ACM SIGCOMM*, 2007.

[12] D. Roselli, J. R. Lorch, and T. E. Anderson. A comparison of file system workloads. In *Proc. of USENIX Annual Technical Conference*, 2000.

[13] 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.

[14] Moritz Steiner, Taoufik En-Najjary, and Ernst W. Biersack. A global view of kad. In *IMC '07: Proceedings of the 7th ACM SIGCOMM conference on Internet measurement*, pages 117–122, New York, NY, USA, 2007. ACM.

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

[16] Mojtaba Torkjazi, Reza Rejaie, and Walter Willinger. Hot today, gone tomorrow: On the migration of myspace users. In *Proceedings of the 2nd ACM SIGCOMM Workshop on Social Networks (WOSN'09)*, 2009.

[17] Xiaoming Wang, Zhongmei Yao, and Dmitri Loguinov. Residual-based measurement of peer and link lifetimes in gnutella networks. In *INFOCOM*, pages 391–399, 2007.

[18] Walter Willinger, David Alderson, and Lun Li. A pragmatic approach to dealing with high-variability in network measurements. In *Internet Measurement Conference*, pages 88–100, 2004.