Resources availability for Peer to Peer systems

Georges Da Costa
Georges.Da-Costa@imag.fr
Laboratoire ID-Imag
Grenoble, France

Corine Marchand
Corine.Marchand@imag.fr

Olivier Richard
Olivier.Richard@imag.fr

Jean-Marc Vincent
Jean-Marc.Vincent@imag.fr

***Abstract*— Nowadays, Peer to Peer systems are largely studied. But in order to evaluate them in a realistic way, a better knowledge of their environments is needed. In this article we focus on the computers availability in these systems. We characterize this availability behind ADSL lines and we link it with the availability of Peer to Peer systems participants. We emphase on the methodology as generalized in other systems such as grids or ad-hoc systems. We finally show how users of ADSL lines are related to Peer to Peer users and we give some examples of the possible practical use of theses results. The results are based on trace datasets obtained over the first five month of 2003 with around 5000 hosts.**

Keywords: Peer to Peer, availability, performance evaluation, model, adsl

I. Introduction

Study of Peer to Peer systems [12] is a recent field of research with a great importance according to technology developments of networks. *The Peer to Peer systems are composed of a large number of participants that cooperate in order to achieve a goal (sharing file, seeking protein forms,. . .). These participants are volunteers and so are highly dynamics and heterogeneous.* Behavior of such systems on a large scale (several million of computers) is still quite poorly known [6] because it is influenced by a great number of factors (in particular : number of connected users, individual user behavior, network architecture and protocols). At the present time, there is a lack of modeling and experiments, which permits their extensive study. Yet the identification and quantification of these factors is fundamental to their understanding. Moreover, the three usual methods of system analysis which are modeling, simulation and emulation, are all based on a clear characterization of the system.

Environment of Peer to Peer systems, featured by all the factors, can be separated in two parts. On the one hand, we consider factors referring to physical architecture, such as characteristics of computers on which the system runs (including CPU speed, memory size, disk capacity), or network (topology, bandwidth, latency,. . .). On the other hand, we gather factors related to the workload of these systems, in particular those concerning specific requests[5], [9].

Unlike classical systems analysis, one should note that Peer to Peer workload is directly related to architecture. For example the amount of available resources induces traffic between computers. So interaction between dynamicity of architecture and workload should be studied very carefully.

The capability of such a system depends on the number of available resources (computers connected through the Peer to Peer network) and their capabilities. To obtain estimation of it, we have to analyze the behavior of individual users, classify and aggregate them.

Based on experiments, the aim of this paper is to propose experimental methods to analyze traces of Peer to Peer architecture evolution, in order to set a model and improving P2P algorithm.

In this article, we describe the availability [1] model of computers taking part in Peer to Peer systems. Indeed, it is difficult to build an a priori model of the time when the users run their computers. There are a multitude of users behaviors ranging from an occasional use to a permanent connexion. On top of that, there is a huge number of users. So the extraction of a complete behavior is complex. Yet, just obtaining this model is a complicated task, as it is not easy to directly obtain information on this population. In addition, using partial data on the availability requires the evaluation of their relevance, as well as their representativeness.

Previous studies[14], [3] have evaluated the availability of computers on Peer to Peer systems. However it is difficult to take into account all the systems. For example [14] studies specifically FastTrack, Gnutella, and Direct-Connect. Then, it is difficult to extract information on the behavior of others systems. The second difficulty is that many systems make it difficult to obtain information on their participants[4].

With this intention, the approach adopted in order to characterize these factors is a trace analysis. This study is based on the French cooperative project of quality evaluation of the ADSL network : *La Grenouille*[8].

The paper is organized as follows. After this introduction, we describe traces issued from experiments on *La Grenouille*. Then we extract global characteristics of these data. Next, we develop in section 4 a model of users profiles based on clustering. By aggregation of users profiles, we get in section 5 the global resource profile. In the section 6, we show an example of generator using the previous results. Finally, we generalize our results from ADSL to Peer to Peer availability.

II. Trace description

We based our study on the data provided by the team *Groupe Clusters Grilles/Grand-Large* (Inria Futurs) within the framework of the ACI-Grid CGP2P. These data are drawn from an attempt of quality evaluation of the French ADSL

	Average	Std deviation
Number of different users per day	4165	203
Number of different users per hour	1402	432
Number of presence days for each user	4.3	2.2
Number of presence hours for each user	34	40

TABLE I
STATISTICS OVER THE FIRST WEEK OF 2003 FOR THE 6852 DIFFERENT COMPUTERS USING *la Grenouille*.

access called *la Grenouille* [8]. The purpose of this project is to provide the *weather of the Net*. For that, the web site of *la Grenouille*[8] publishes in real time the performances of various ADSL providers.

This project is independent - meaning it is not related to any of the various operators providing ADSL access - and provides unbiased information. Moreover, the criteria used in the calculation of offers qualities are known, which reinforces the objectivity of this site. Finally this project has interesting features for the users. It allows users to compare the overall quality of the different offers helping them know the punctual problems.

Indeed, it is an open source project (its source code is available), it is transparent for the user and consumes few resources. These characteristics lead to think that a great number of people with various interests can use this system, thus guaranteeing its representativeness. Finally, this system is largely used in France. The participants are spread homogeneously over the French territory, and each providers has a stray amount of users. This will be even more true in a near future as some french providers wish to put this systems on their official connexion kit. It would allow them to obtain information on the state and quality of the users with problems. In this case, this system would become perfectly representative.

This study is intended for large wide-spread Peer to Peer system used by volunteer, not by organization for using their idle power, such as in[2]. In [2], the authors study the availability of 50000 corporate computers.

Our study was restricted to personal computers connected to the Internet by the ADSL. Then we chose not to take into account computers such as servers of providers or computing centers.

The participants, composed of ADSL users who had chosen to help the project of *la Grenouille*, are divided among two categories : private individuals and access providers. The basis of this project is that each user takes measurements of quality of its access to the Internet. The obtained data are then centralized on the servers of *la Grenouille* which then evaluate the quality of the ADSL offers and finally print results on its web site.

These data are divided into two files : one file containing the category of each participant, and one file containing the results of the measures taken by each participant.

Furthermore each one of these measurements is represented by various fields like : the identifier of the trace, the date,

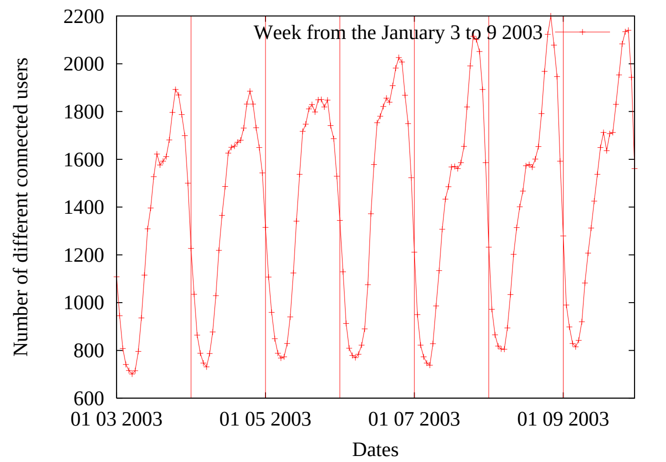

Fig. 1. Evolution of the number of different computers connected on the Internet over one week

the hour of the measurement and an single identifier of the computer. For example :

```
Id      Date         Hour       Computer
4428    2003-01-09   12:00:00   25363
```

These measurements were beginning in 2000 and are still in progress. As the measuring unit considered is the half-hour, it is thus not possible to obtain information of finer granularity than half-hour. The number of the measurement (the single identifier) is set by the server of *la Grenouille*, contrary to the date and the hour of measurement, which are the local time of the computer that takes the measurement.

Also, in order to be able to process these data and compare them, the assumption of synchronization of the clocks is necessary. However the detection of the desynchronized clocks is possible with a precision of about half an hour which is in any case the current maximum precision of the system. Thus, this synchronization is possible considering two facts. The first one is that we can suppose that most computers have a clock precise at a quarter-hour. In this case, as the granularity of our measurements is twice this time, it will not affect it. The second facts is that the data received are received with the one of the other computers. In this case, it becomes possible to compare them in order to change any absurd entry. For the data we manipulated, there were a lot of false occurrences of 00:00 for example.

Besides, it is not impossible that certain measurements are lost because they are obtained locally before being sent to the server.

The total size of all these data which we are manipulating is about ten Gigabytes.

III. GLOBAL ANALYSIS OF THE DATA SET

Considering the large amount of data, we have chosen to subdivide the data per year and we focused the present study over the year 2003. The data before 2003 are not as precise due to a rewrite of the client at the end of 2002[1]. Most of

[1]The pre-2003 data confirm our following results.

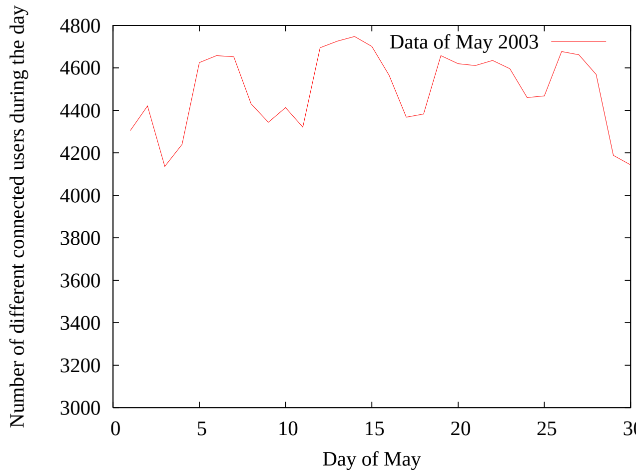

Fig. 2. Evolution of the number of different computers connected over one month

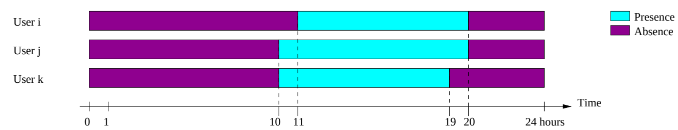

Fig. 3. Users i, j and k have a close behavior over one day. i and k are at a distance one by construction

the following has been verified on other dates than the one displayed on the graph. They have been verified on at last the entire months of April and May 2003.

Moreover to ensure simplification, a preprocessing step was applied to the raw data in order to eliminate measurements referring to the access provider.

In addition, within the framework of Peer to Peer systems, after having established our criteria of interest, some measurements can be eliminated. For example a user being connected only once during the year will be regarded as non-representative. It must be someone who tried the system *la Grenouille* and who is no more participating in it.

After this first phase of elimination, it is possible to consider various approaches. The statistical analysis can be carried out from a user point of view (a number of different users per day, per hour...), and from a temporal point of view (for example, the number of days of presence for each user).

A synthesis of some empirical results is presented on Table I. This first analysis is aimed to check the quality of the data.

First, we have to determine that a large number of users are using the system on a regular basis. Thus it is important to verify that most of the users are present on a regular basis. This table shows that people are using the system quite often, as they use it more than 4 days on average. In this case it could be only power users. But as the standard deviation of the number of presence hour is 40 (compared to an average value of 34), it appears that there are a lot of different behaviors.

The number of simultaneous people underlines the large number of different behaviors. There is never more than 61% of the people that are connected at the same time.

Finally, we can verify that the users have a globally constant use of the system as the standard deviation of the number of different users per day is small (less than 5% of the value).

In order to reduce the field of investigation, we were interested in the evolution over the time of the number of connections (different connected computers). Figure 1 shows a daily periodicity of this number of connections. As the day is the time basis of most of our behavior it seems that the daily periodicity can be first studied. There are enough information in a day (5000 different users connected during one day on a sample of about 7000 users) and thus the sample considered remains representative. Table II shows that for each day of the first week of 2003, the number of connected computers for each day remains high enough compared to the total number of computers to be representative.

This daily periodicity is also found in greater period of time. Thus it is visible on Figure 2 that there is also a periodicity of about a week, which is less marked than the daily one. Further experiment showed us that the symmetry followed the living rhythm such as holidays, or public holidays.

The choice to restrict our study at one day is justified by the complexity of the problem. This simplification will certainly imply a loss of information. Also after the exploitation of the data over one day, a general study over longer periods will be necessary to emphasize other relevant information.

IV. Users profiles

To treat such a large number of data, the scientific difficulty consists in identifying the significant data. In our approach we defined as being significant the hours and connection times of each users. Thus, a standard behavior is defined by a large number of users adopting a globally similar attitude when it comes to their hours and connect times during the day. Actually, considering the sample size, it is necessary to gather the users in a first class repartition. From this point on, each user will be represented by a vector of Boolean indicating his connection on each of the 24 hours of the considered day. Thus by regrouping the users which have the same behavior, the first classification reduce the amount of data to process.

However this first classification remains too indistinct. Effectively the 4373 users connected in the day of January 7 2003 are represented by 1710 different classes. In order to gather these various behaviors in classes there are several methods. We started by carrying out a principal component analysis [11] to obtain the first classification. This method is particularly used when it is necessary to obtaining a vector base of a space. But this approach did not enable us to obtain relevant results.

Consequently, the question is to find according to which criterion we can fusion several classes ? Our method then consists in calculating the Hamming distances between each class to obtain a finer classification joining together similar behaviors. The Hamming distance [15] is a simple method to compare two lists. It is defined by the minimal number of elementary operations needed to transform a list into the other. These basic operations are the modification, the suppression

Date (2003)	01 03	01 04	01 05	01 06	01 07	01 08	01 09
Number of users	3850	3985	3994	4229	4373	4384	4343

TABLE II

NUMBER OF DIFFERENT PERSONAL COMPUTERS CONNECTED PER DAY ON A TOTAL FROM 6852 DIFFERENT COMPUTERS BEING CONNECTED OVER THE FIRST WEEK OF 2003

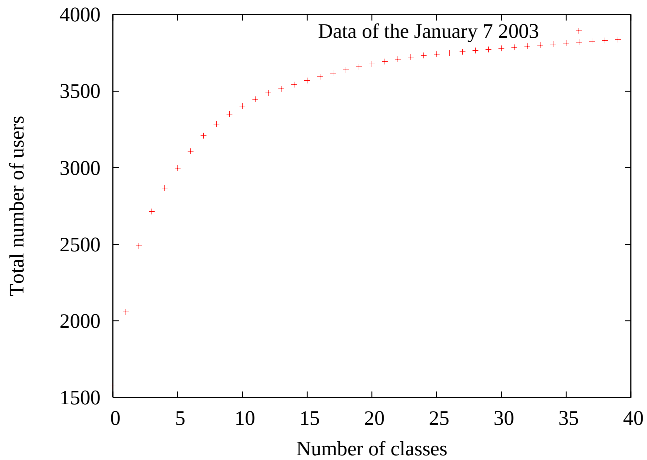

Fig. 4. Size of the sum of the largest classes in a number of computers

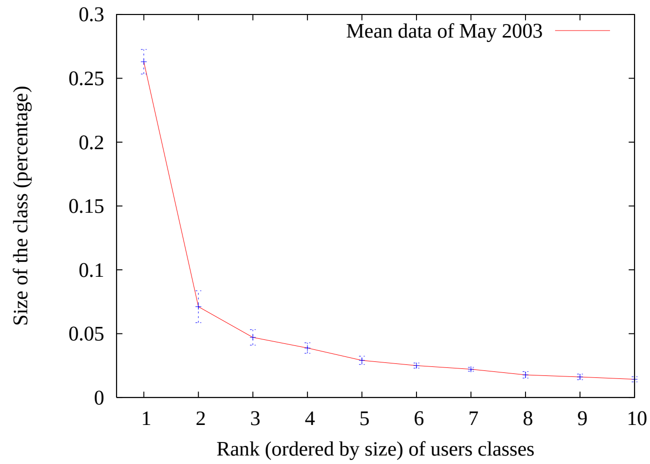

Fig. 5. Comparison of the clusters sizes for one month

and the addition of an element. Moreover the weights can be chosen for each one of these operations.

In our context, in order to take into account at the same time the hour of connection and its duration, we have chosen to emphasis the removal and the addition of an element opposed to the operation of modification. Thus, users i , j and k represented on Figure 3 are all at the same distance (within the meaning of Hamming) of each other. The elementary operation of modification has a weight of 1, and the elementary operations of remove and append have a weight of $1/2$.

Once the Hamming distances obtained between the various classes of users, several methods of joining them is possible [10]. One can for example use a neighborhood graph and, for a given user, form a class with its direct neighbors, the neighbors of the neighbors, and so on until a given depth. For example, by regarding a Hamming distance of 1 between neighbors and with fixing the depth to 2, the new classification extracted has 506 different classes.

So, there are several ways to construct the classes. We tested the different approaches and verified that the resulting size laws were the same. Besides the classes were almost the same. So we focused our experiment with the *fixed distance*.

As Figure 4 shows it , the majority of the users are integrated into a small number of classes. This evaluation of the class sizes was reproduced over several days in order to obtain the characteristics of these classes. One can see on Figure 5 that the size distribution of these classes is nearly constant over May 2003 with respect to time with a coefficient of variation lower than 0.15.

On these Figures (4 and 5) illustrate the distribution of the classes in two categories. On one side, a low number of classes including most of the behaviors, on the other side a great number of behaviors (approximately a quarter) rather isolated whose distribution remains to be discovered. The class sizes are similar on each days and, except for the representativeness, the extracted classes are identical one day on the other.

It is to be precised that the size of this class follows a Zipf law.

V. RESOURCE PROFILES

These measurements also enable us to gather information on the use of the resources, in particular on its evolution during the time in term of connected computers number.

A characteristic phenomenon confirms our intuition, namely the raise of the resources available during one day (see on Figure 1). This increase is split into two phases separated by a plateau, which corresponds to the period ranging between 12:00 am and 2:00 pm.

The previous classes analysis allows to provide an explanation of the plateau apparition. A large number of computers are gathered in two sets of classes : the classes for which the activity is centered on the morning, and the classes for which the activity is centered on the afternoon. Thus as the classes present on the afternoon gathered more people than the one on the morning, the peak of useness is on the afternoon. The plateau is when the computers connected on the morning disconnect at the same rate as the connection of the computers in the afternoon. The first part (composed by users connected only between 04:00 am to 2:00 pm) is composed by the 14 of the 30 largest classes, and the second one (from 12:00 to 1:00 am) is composed by 12 out of 30.

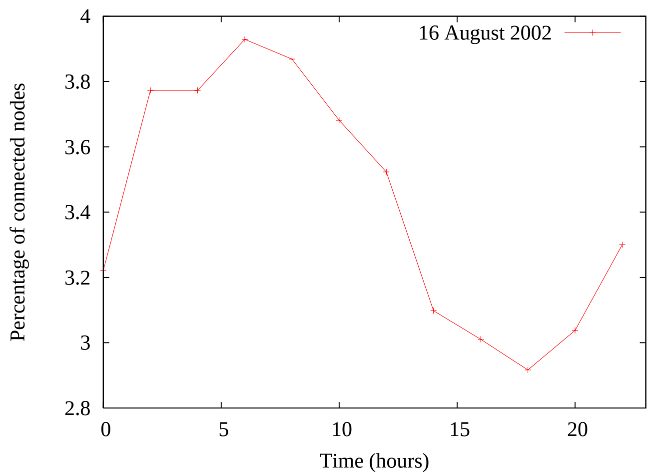

Fig. 6. Percentage of the Gnutella users present the August 16 2002

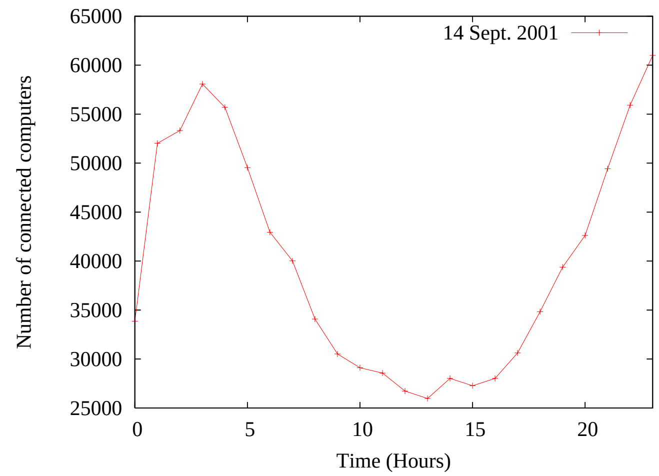

Fig. 7. Number of connected computers on the Kazaa system the September 14 2001

VI. Availability profile generator

In the following, we will provide an example of generator which would lead to generate the availability profile of computers which could be used in a simulator or an emulator. This will be part of *Rig*, a realistic inputs generator for Peer to Peer systems evaluation.

For a set of n nodes, the daily profile will be generated as following :

For each node, using the Zipf Law, choose the class it will belong to.

If it is one of the 100 first class, then take the representative, and add some noise (staying at least at a Hamming distance of 1 of the representative). Else, generate randomly a behavior with a mean availability of 6 hours.

We choose to use the 100 first classes as after this value all the classes seems to be evenly distributed. We measured the mean availability of the remaining classes in order to be able to generate them.

VII. Links between the Peer to Peer and ADSL availability

We have provided an analysis of the ADSL computers behavior. There is a link between these users and the participants of the Peer to Peer systems. Indeed it is currently a connection (like the cable) which offer the qualities needed to achieve a good participation in such systems. The bandwidth is sufficient and the fee are monthly based (and not per byte based).

To emphase the correlation between the two characteristics, we have used two other sources of data. The first one is a Gnutella[7] crawler we have implemented that measure the number of clients reachable in the Gnutella system. Figure 6 shows that the number of clients follows a law that is near to the ADSL one.

The second one is from a study[14] of the Kazaa network. Figure 7 suggests the same pattern of use.

Other studies such as [13], [3] show such a behavior. The deep difference between these studies and ours is that they are linked on one particular system. Moreover, if the system studied becomes obsolete, the methodology becomes too.

VIII. Conclusion

We described and implemented a methodology to characterize the availability of computer behind the ADSL network. Then we have provided a characterization of this availability. Finally we have underlined the link between Peer to Peer and ADSL availability. The presented results have been verified during several months (from January to May), proving the representativeness of the results.

This work has been done thought the useness of several giga bytes of data provided by the *la Grenouille* project.

This study shows the large variability of the resources. Thus systems that keep information on the participant computers, need to choose between : having to wait in order to have the information, or using enough redundancies to always have one participant who holds a piece of information.

It leads to an estimation of the length of connection that can help to tune the peer to peer systems. The classes permit to differentiate the roles of the users in function of their connection pattern.

IX. Future works

In this article, we rapidly passed over the scaling problem due to the lack of space. Yet it is necessary to verify that the laws which have been obtained with a certain size of problem are still adapted at an other size.

In a more generic fashion, the methodology followed in this article can be adapted to many fields. In particular it can be applied to systems which have similar constraints of connections volatility (for example, wireless-networks).

However it remains to extract a global model from these standard behaviors and their statistical laws. A meticulous validation will be necessary because the objective of such a model is to implement a generator. For that we envisage to use two systems together. The first consists of a daily generator.

It will directly use the results extracted from this study. The second one will consist of a generator of days to provide links between the days. Indeed the coherence of a users behavior is proven only over one period of length higher than one day. We envisage to use a finite-state machine to describe the evolution during the time of this behavior.

X. Acknowledgement

We wish to thank the *la Grenouille* project for its data.

References

[1] Ranjita Bhagwan, Stefan Savage, and Geoffrey M. Voelker. Understanding availability. IPTPS03, 2003.

[2] William J. Bolosky, John R. Douceur, David Ely, and Marvin Theimer. Feasibility of a serverless distributed file system deployed on an existing set of desktop pcs. In *SIGMETRICS '00: Proceedings of the 2000 ACM SIGMETRICS international conference on Measurement and modeling of computer systems*, pages 34–43, New York, NY, USA, 2000. ACM Press.

[3] J. Chu, K. Labonte, and B. Levine. Availability and locality measurements of peer-to-peer file systems. ITCom: Scalability and Traffic Control in IP Networks., 2002.

[4] Ian Clarke, Oskar Sandberg, Brandon Wiley, and Theodore W. Hong. Freenet: A distributed anonymous information storage and retrieval system. *Lecture Notes in Computer Science*, 2009:46–62, 2001.

[5] G. Da-Costa and O. Richard. Impact of realistic workload in peer-to-peer systems, a case study : Freenet. ISPDC Proceedings, 2002.

[6] Neil Daswani, Hector Garcia-Molina, and Beverly Yang. Open problems in data-sharing peer-to-peer systems. In *ICDT '03: Proceedings of the 9th International Conference on Database Theory*, pages 1–15. Springer-Verlag, 2003.

[7] Gnutella, http://gnutella.wego.com/.

[8] La grenouille, http://www.grenouille.com.

[9] Krishna P. Gummadi, Richard J. Dunn, Stefan Saroiu, Steven D. Gribble, Henry M. Levy, and John Zahorjan. Measurement, modeling and analysis of a peer-to-peer file-sharing workload. 19th ACM Symposium of Operating Systems Principles (SOSP), October 2003.

[10] R. Jain. *The Art Of Computer Systems Performance Analysis*. John Wiley and Son, 1991.

[11] M. Piron L. Lebart, A. Morineau. *Statistique exploratoire multidimensionnelle*. Dunod, 1995.

[12] Andy Oram. *Peer-to-Peer : Harnessing the Power of Disruptive Technologies*. O'Reilly, 2001.

[13] Stefan Saroiu, P. Krishna Gummadi, and Steven Gribble. A measurement study of peer-to-peer file sharing systems. In *SPIE Multimedia Computing and Networking (MMCN2002)*, 2002.

[14] Shubho Sen and Jia Wang. Analyzing peer-to-peer traffic across large networks. *ACM/IEEE Transactions on Networking*, 12(2), Avril 2004.

[15] A. Tanenbaum. *Réseau*. Dunod, 1997.