

ETOMIC Advanced Network Monitoring System for Future Internet Experimentation

István Csabai¹, Attila Fekete¹, Péter Hága¹, Béla Hullár¹, Gábor Kurucz¹, Sándor Laki¹, Péter Mátray¹, József Stéger¹, Gábor Vattay¹, Felix Espina², Santiago Garcia-Jimenez², Mikel Izal², Eduardo Magaña², Daniel Morató², Javier Aracil³, Francisco Gómez³, Ivan Gonzalez³, Sergio López-Buedo³, Victor Moreno³ and Javier Ramos³

¹ Eötvös Loránd Univeristy, Budapest, Hungary

{csabai, fekete, haga, hullar, kurucz, laki, matray, steger, vattay}@etomic.org

² Universidad Pública de Navarra, Pamplona, Spain

{felix.espina, santiago.garcia, mikel.izal, eduardo.magana, daniel.morato}@etomic.org

³ Universidad Autónoma de Madrid, Madrid, Spain

{javier.aracil, francisco.gomez, ivan.gonzalez, sergio.lopez-buedo, victor.moreno, javier.ramos}@etomic.org

Abstract. ETOMIC is a network traffic measurement platform with high precision GPS-synchronized monitoring nodes. The infrastructure is publicly available to the network research community, supporting advanced experimental techniques by providing high precision hardware equipments and a Central Management System. Researchers can deploy their own active measurement codes to perform experiments on the public Internet. Recently, the functionalities of the original system were significantly extended and new generation measurement nodes were deployed. The system now also includes well structured data repositories to archive and share raw and evaluated data. These features make ETOMIC as one of the experimental facilities that support the design, development and validation of novel experimental techniques for the future Internet. In this paper we focus on the improved capabilities of the management system, the recent extensions of the node architecture and the accompanying database solutions.

Key words: ETOMIC, network measurement infrastructure, active measurements

1 Introduction

Internet is existing since the late 60s and became a global network in the early 80s. By that time the general design and most of the basic protocols have been laid down. Since then its scale and complexity has become orders of magnitudes larger concerning both the number of nodes and the amount of traffic flowing through. Although smaller modifications have happened and the hardware behind the Internet has changed significantly, the basic structure has remained the

same. Recently, several international efforts try to re-design the Internet from clean slate. It is expected that a more optimal design will solve the most pressing problems we currently face. To enable the design of a better system it is essential to understand the details of the network and traffic dynamics. For this purpose we need to create measurement tools and gather diverse measurement data.

The European Traffic Observatory Measurement InfrastruCture (ETOMIC)¹ was launched in 2004 [1, 2, 3] with the need to supply the network research community with various experimental data. It is targeted to provide the scientific community with an Internet measurement platform that is fully open and reconfigurable, extremely accurate and GPS-synchronized. The ETOMIC system has been designed to allow researchers to perform any kind of active network measurement. The users are provided with a web-based graphical user interface for the definition of the experiments to run. Researchers may either choose from a number of built-in scripts that cover the most popular measurement techniques, like traceroute or packet-pair experiments, or they can provide their own code for the experiments. To avoid conflicts in resource utilization each measurement has to be scheduled to exclusively reserve node resources for its execution. The node reservations are performed through the web-based user interface. The ETOMIC management kernel takes care of the software upload and experiment execution in a fully automated fashion.

After the successful duty of the measurement nodes since 2004 the renewal of the system components was necessary. In the OneLab project [4] we have extended the capabilities of the measurement hardware to match the current technologies and to incorporate the software evolution of the last years that are important from the perspective of network measurements. The ETOMIC infrastructure now provides two ways of collecting experimental data. One possibility is when the researcher reserves and configures the measurement nodes and sets the parameters of the experiment through the Central Management System. In this case, besides the original ETOMIC nodes, newly deployed enhanced measurement boxes can also be used for experimentation. To meet the requirements of high precision measurements the nodes are equipped with a DAG card (for the original nodes) or an ARGOS card (for the new generation nodes) to provide nanosecond-scale timestamping of network packets. Besides these nodes a third type of hardware component was also introduced, which is called Advance Probing Equipment (APE). APE is a low cost hardware solution developed to serve as a measurement agent for user applications: it provides a web service interface to conduct experiments. This approach enables autonomic software components to automatically collect relevant network data from the ETOMIC system they rely on for their operation. As a consequence of a development in the system kernel the nodes of the PlanetLab platform [5] can also be used as measurement nodes by the ETOMIC system. The goal of this integration was to enable the federated usage of the high precision ETOMIC nodes and the numerous PlanetLab nodes.

To make it easier to handle and archive the huge amount of data collected by the ETOMIC platform we have created data repositories. There are two different

¹ ETOMIC was awarded with the Best Testbed Award at TRIDENTCOM 2005.

interfaces for these data archives. The periodic measurements web interface can be used to poll automatically collected measurement data through pre-defined queries. Currently this interface provides one-way delay measurements, NTP and GPS measurements and several different types of traceroute experiments. As another approach, the Network Measurement Virtual Observatory (*nmVO*) [6] provides standard SQL database access to the user community. The nmVO provides a graphical user interface and a web service interface for accessing raw and evaluated measurement data.

The paper is organized in the following way: Section 2 explains the key features of the system architecture. In Section 3 and in Section 4 we introduce the management kernel that provides the central control of the system and the hardware architecture of the measurement nodes. Section 5 discusses the different data repositories and their interfaces. Finally, in Section 6 the conclusions close the paper.

2 System architecture

The ETOMIC infrastructure is constituted of high precision measurement equipment modules hosted by European universities, research institutes and company laboratories. The clocks of the measurement nodes are synchronized via GPS signals, which allow not only packet round-trip time estimation, but also precise one-way delay measurements. The ETOMIC platform is very flexible, since researchers can develop and run any kind of active experiments.

A Central Management System (CMS) is in charge of system control, comprising not only the scheduling and execution of measurements experiments, but also system monitoring and configuration. The main software component of the CMS is the *management kernel* which is running on a dedicated server computer. The kernel is responsible for scheduling tasks, deployment of user software to the measurement nodes, node configuration, experiment execution and the collection of measurement results from the nodes. The CMS provides a web-based graphical user interface where the researchers can configure the system and reserve system resources for their measurements. An internal database is attached to the kernel where the system and user level management information are stored. The results of a finished measurement are also collected and stored in this database until the user downloads them. The system components are depicted in Figure 1.

As an important add-on for the original ETOMIC system the architecture has been extended with large capacity data repositories that are publicly available. The system provides several different interfaces for these repositories through which the users can reach the collected datasets. The interfaces allow the users to run intelligent queries in order to filter and process the raw data on the server side. In these repositories the results of traceroute, one-way delay, queuing delay tomography, available bandwidth, IP geolocalization, router interface clustering and NTP-GPS measurements are stored.




Fig. 1. ETOMIC system architecture.

3 Management kernel

The ETOMIC management kernel constitutes the core of the Central Management System (CMS). It is in charge of user management, experiment scheduling and keeping the corresponding results in the temporary data storage. In the following we briefly overview the tasks performed by the CMS.

Scheduling and calendar maintenance In order to isolate the different measurements and to schedule experiments and maintenance tasks a calendar is used for each measurement node. The researcher is expected to book measurement nodes for a certain time interval and to upload the applications necessary for the measurement. The web interface is in charge of checking that the time line for the experiment does not collide with any other previously registered measurement. In case of successful resource reservation the CMS inserts the new experiment information into its internal database. Previously, the only way to create periodic experiments was to define their scheduling one-by-one. Due to a recent development in the kernel scheduler the CMS allows the definition of customized periodic experiments with an inter-experiment time interval and a repetition count. This feature enables to program for example daily or hourly experiments with the ease of defining only one experiment.

Experiment management The management kernel is continuously checking the internal database for new measurement requests. When there are no scheduled experiments the measurement nodes are totally stateless for failsafe reasons. The communication is always initiated by the server via SSL based secure connection to command the nodes. Once a new experiment has been defined and the deadline for execution approaches, the management kernel performs the following tasks (that are internal to the system and transparent to the researchers).

- **Software upload and measurement node configuration.** The measurement applications and accompanying files are uploaded and the nodes are configured right before the experiment starting time. The measurement nodes

are programmed with the starting and ending time and the parameter sets for each of the executables that are going to be run for the experiment.

- **Experiment execution.** During this task the management kernel is on standby until the end of the experiment. It only has to ensure that no multiple experiments are using the same measurement nodes in the same time. The aim is to completely isolate the measurement nodes so that the high-precision measurement hardware can be exclusively devoted to a single experiment.
- **Results download.** Once the experiment is finished, the management kernel downloads the resulting data files from each measurement node to the temporary data storage. Since the download time is highly dependent on the network connectivity and the size of the output data files we incorporated resume capability into the file transfer. In this way the data transfer can be pended and resumed according to other scheduled experiments in order to avoid interference between data transfer and the experiment execution. With these functionalities the CMS is able to manage even very large datasets. After successful data download all the nodes are set back into the initial state and the remaining user files are deleted.

Periodic measurements as maintenance tasks Due to the upgrade of the management kernel not only the users are able to define periodic experiments but also the CMS itself can carry out periodic measurements. A new kind of low priority task has been introduced into the kernel scheduler. These periodic measurements are planned as management tasks into empty slots of the calendars of the measurement nodes and executed only if the nodes are idle. If a node has become reserved for that time slot by a user experiment then the management task is canceled and automatically re-scheduled for a later time period and the user's experiment is executed. For the end-users these measurements are not visible and they do not affect the availability of the nodes.

3.1 Web interface to CMS

ETOMIC provides an interface for researchers and administrators which fulfills the different requirements they may have. The former require capabilities to define new experiments, to reserve the measurement nodes and to download results from the CMS. The latter require functionalities to manage users, measurement nodes, software and experiments. An internal database is used to store all the necessary information to run the experiments. The stored information includes the applied softwares, external data files, the experiment results, the experiment status and the measurement node status. The strong usability requirements imposed by the researchers and administrators represented a significant challenge for the interface design. Solutions based on providing console access (`ssh` or `telnet`) to the nodes during the experiment were not deemed adequate, in order not to burden the researcher with launching the software at each node and with scheduling the tasks during the experiment. Instead, users are provided with a graphical interface for setting up the experiment beforehand. Then, the management kernel is in charge of experiment execution. The interface is based

on the ubiquitous web service. The aim is to facilitate the definition of the experiments as much as possible. This targeted simplicity and ease of use posed more stringent challenges on the researchers interface.

User interface The main user interface used by the researchers is described in depth in [2] and [3]. Users can find manuals, the programming API and example codes here. Next we highlight the main functionalities and explain only the new features developed since the release of the original ETOMIC infrastructure.

- **Adding a new program**, executable files or source code of the measurement.
- **Uploading data files** if it is needed in addition to the executable files.
- **Creating an experiment-bundle** to define the experiment with scheduling the start and end times of the measurement.
- **Booking ETOMIC time** to reserve the measurement nodes and the certain time frame for the experiment.
- **Obtaining results**, downloading the resulting data files from the temporary storage to the user’s own computer.
- **Periodic experiments** can be defined with a repetition period and the system with automatically schedule as many experiments as requested and allowed by the user privileges. Note that these periodic experiments run with normal priority, their time and node reservation is ensured by the CMS in contrast to system level periodic measurement described in Section 3.
- **Sharing files**: the users can share their executable files, data files or results directly with other ETOMIC users. The users are also welcome to publish their results in the open repository or to store their raw experimental data in the Network Measurement Virtual Observatory (see Section 5.2).
- **Web service interface**: a REST web service allows to download experiment results from scripts that the researchers may want to develop.

Administrator interface Different administrator types are in charge of the ETOMIC system. The CMS and the measurement nodes are centrally administrated, but there are local administrators per measurement node for the physical maintenance. The central administrative tasks include the user management, to open new accounts, change user privileges and maintain the internal databases. Finally, the ETOMIC superuser has privileges to access and modify low-level system information and to abort any experiment programmed or in progress. All these features are provided through a user-friendly web interface.

3.2 Integration of Planetlab’s nodes

PlanetLab [5] is a global platform for supporting the development of new network services. This platform is also used for network experiments. The PlanetLab nodes are accessed interactively via remote shell. This access method enables the CMS to use the PlanetLab nodes as its own nodes. Although the main hardware capabilities of the PlanetLab and ETOMIC nodes significantly differ, the large number of PlanetLab nodes makes them very attractive to the user community. The capabilities of PlanetLab are not described in this paper, here we

only note that the PlanetLab nodes are usually up-to-date server PCs without any hardware components specialized for network measurements. The slice based management of PlanetLab nodes allows multiple users to run experiments simultaneously in the same remote node at the same time, while the CMS takes care of the unique resource allocation. In spite of the basic differences of PlanetLab and ETOMIC the federated usage of the high precision ETOMIC nodes and the numerous PlanetLab nodes could lead to new ways of experimentation.

The software installed on ETOMIC nodes has been adapted to make the joint usage possible, using a slice of PlanetLab that is automatically renewed by the CMS. This makes the whole range of ETOMIC and PlanetLab remote nodes available through the ETOMIC web interface. The most important challenge for the integration was the synchronization of the clocks in nodes from both platforms as they use different reference signals with highly different precisions. Using the ETOMIC interface, a software that is transparent to the end user takes into account the time difference between the nodes in order to allow synchronized launch of experiments at all the nodes.

4 The measurement nodes

In this section the hardware and software components of the ETOMIC measurement nodes are presented. The nodes can be divided into two groups based on their hardware architecture. The ones that are built on server PC architecture are called *ETOM*. These nodes are accessible via the web-based graphical interface presented in Section 3.1. The ones that are based on a lightweight programmable board are called *APE*. The APE nodes are accessible via a web service interface. GPS receivers are connected to all types of measurement hardware to provide the precise time synchronization between the nodes and to provide the reference clock for the measurement cards.

Each of ETOMIC's hardware solutions is high-precision, due to the incorporated precision equipment that are specifically designed to transmit packet trains with strict timing, in the range of nanoseconds. Each measurement node is provided by two network interfaces: a standard network interface card for management purposes (maintenance, software upload, data download) and an additional precision card for the probe traffic.

4.1 ETOMs with DAG card

The first version of the measurement nodes are based on Intel S875WP1E server PC architecture with Debian Linux operating system with enhanced kernel capabilities for low level network access without root privileges. For the network monitoring interface an Endace DAG 3.6GE is used, which is specifically designed for active measurements. Such cards do not use interrupts to signal packet arrivals to the kernel, and thus packets can be captured at gigabit speeds. Shared memory is used as means to relay packets to the analysis program running in

user space, in such a way that interrupts are avoided. Furthermore, the GPS reference signal is used to timestamp the incoming packets with high resolution. To program the DAG cards a C language API and some user space applications are provided, e.g. a packet pattern sender and a packet capturer.

The GPS reference is based on a Garmin GPS 35 HVS. It is a water-resistant GPS receiver that is used to synchronize the measurement nodes. Specifically, the GPS provides a pulse per second signal directly to the DAG card. The resulting accuracy is 60 nanoseconds in the packet timestamps and inter frame generation intervals. A signal converter has been designed in order to bridge the mismatch between GPS receiver and PC ports. The range of the transformed signal allows large distances between the GPS receiver itself and the PC which enables good sky visibility for the GPS receiver unit.

4.2 ETOMs with ARGOS card

The new generation ETOMs are based on a HP ProLiant ML370 server PC architecture with Ubuntu Linux. The quad-core server processor provides sufficient computation power to carry out complex measurements and online data analysis. Two heavy-duty disks are responsible for the safe and fast data storage during the measurements. These nodes are equipped with ARGOS measurement cards that are based on the netFGPA platform of Digilent Inc. ARGOS is capable of timestamping IP packets with nanosecond precision. In contrast to the DAG cards, it can act as a standard network interface card. Practically, this solution improves the development of probing applications, since users can program the cards through the standard socket interface and query the packet timestamps with the `libpcap` library. A few libraries and kernel modules were modified to support `sk_buf` structures with nanosecond resolution, which in combination with the ARGOS driver, makes it possible to read data directly from either the socket interface or `libpcap`.

In order to provide precise time synchronization a U-Blox LEA-4T high performance, high precision timing GPS is used. The LEA-4T features Time Mode function that enables GPS timing with only one visible satellite and eliminates timing errors which otherwise results from positioning errors. The precise geographic location of the GPS device should be provided once in advance.

4.3 APE lightweight measurement nodes

Besides the development of the ETOMs a novel low-cost measurement equipment was designed which is called Active Probing Equipment (APE). The main feature of the APE nodes is that they provide measurement services which can be remotely called by user applications in an online fashion, without time slot reservation. The APE is built on a development board with Blackfin processor. The board is manufactured by Analog Devices Inc. and has a number of different interfaces for hosting auxiliary hardware components that are responsible for specific network measurements tasks. For instance, a flash memory card was integrated onto the board to enable the temporary storage of measurement data.

Table 1. Available measurement nodes in the ETOMIC system.

	ETOM w DAG	ETOM w ARGOS	APE	PlanetLab
platform	Intel server PC	HP server PC	Blackfin board	variable
timestamping accuracy	60 ns	10 ns	100 ns	$\sim 10\mu s$
time synchronization	yes	yes	yes	no
GPS receiver	Garmin 35HVS	U-Blox LEA-4T	U-Blox LEA-4T	–
number of deployed nodes	18	20	20	~ 300
user interface	web GUI	web GUI	web services	web GUI

The APE nodes are installed with uClinux, a special Linux operating system developed for embedded systems. The devices can be instructed via a web service interface, which opens simple ways for researchers to develop their applications using the network measurement services of the APE nodes. The following APE services are already implemented, and can be called (similarly to remote procedure calls): ping, traceroute, an arbitrary packet pattern sender and a capture tool.

5 Accessing measurement data

For network measurements the collected raw data is traditionally stored in files in some standard (like traceroute dump, tcpdump) or custom formats. The files are then processed according to the research questions to be answered. Detailed analysis of complex networks requires large statistical samples. This requirement leads to substantial data size in case of measurements in high bandwidth networks, even if just a few parameters of the packets are recorded (like IP addresses, arrival time, protocol, size or delay). Practically, measurements can produce dozens of megabytes at each monitoring node that sums up to hundreds of megabytes or even terabytes in multi-node experiments. Keeping only the results of the data analysis and discarding the raw data itself is not a good way to solve the data handling issues. Since measurement data gathered today cannot be reproduced in the future, it is preferable to store the original datasets to allow further re-analysis (applying the various different statistical methods to be developed in the coming years), and to support the study of the long-term evolution of the network.

For these purposes we have created data repositories to store measurement data collected by means of the ETOMIC system. There are two different interfaces to reach raw and aggregated measurement data. The periodic measurements web interface and a web service interface can be used to poll data collected from automatic measurements in the system. The users can choose from

pre-defined queries by selecting a given type of measurement, a set of nodes on which the measurement was conducted and a time frame.

We also present another interface which is called Network Measurement Virtual Observatory (*nmVO*). The nmVO approach aims at providing efficient and flexible access to raw experimental data and server side analysis tools [6]. It is also capable of integrating various databases in a common framework that provides full SQL functionality. Researchers may reach the nmVO through a web GUI or a web service interface, where the users can automatically poll data through their own client programs.

5.1 Web interface for periodic measurements

To collect long term traces we have created periodic measurement tasks in order to run specific measurements when the nodes are idle. All the data from the periodic measurements is publicly available through an easy-to-use web interface, where the user can choose from pre-defined queries. For each measurement the user can download filtered data specifying only some pairs of nodes and a period of time. Once the selection is done the user can choose between downloading the data in a compressed text format or viewing a plot for those data provided by the server. In the following we list the types of periodic measurements for which the collected data is available through the ETOMIC website's **Database / Open Repository** menu [1].

NTP-GPS: NTP and GPS synchronization information is collected several times a day from every node. These are measurements collect information from the GPS devices connected to each monitoring node and from the NTP daemon running on the nodes.

One-way delay: Every possible pair of ETOMIC node takes part in periodic one way delay measurements. These measurements are done using the high precision measurement cards and GPS synchronized timestamping.

Traceroute: Traceroute measurements between every pair of ETOMIC nodes are collected. The measurements are done several times a day using standard Ethernet cards.

Paris traceroute: **Paris traceroute** [7] measurements are also conducted between every pair of ETOMIC nodes. The probe packets are designed to always use the same path even if load balancing is being applied in the network. Three types of packets (ICMP, UDP and TCP) are used in order to evaluate different behaviors of the network depending on the transport protocol. The measurements are run several times a day using the Ethernet cards.

5.2 Network Measurement Virtual Observatory

The Network Measurement Virtual Observatory concept, presented in [6] is an approach to efficiently store and share research data. Beyond the simple data

collecting and archiving functions it aims at providing easy-to-use analysis tools via both human and machine readable interfaces. One of the main features of the nmVO is that it provides SQL access for the databases that are integrated under its framework, thus the users can edit and run their customized queries through either the web-based SQL interface or the web services interface. The main advantage of this solution is that the researchers can filter out the relevant information from the huge archives using server side processing. Hence, only the necessary datasets and results have to be downloaded from the server.

To sketch the nmVO principle through a possible application, consider a scenario in peer-to-peer overlay networks where management information is needed to optimize the routing between the peers. It would be unthinkable to use gzipped files for such real-time evaluation. On the contrary, the scenario is feasible if one turns to the nmVO to get the typical loss rate, the average delay on certain routes or the shortest path between the peers. This means that beyond the data itself, analysis tools are also needed to perform such data filtering and transformation queries efficiently. The recent efforts on enabling database engines to run complex user code and define new data types (e.g. MS SQL Server CLR integration) makes this task easier. Using these stored procedures we can move the typical filtering and pre-processing tasks - like getting slices or aggregates of data - to server side. Since end-users can call these functions remotely they can reach the evaluated results with less code and they have to fetch much smaller amount of data from the archive. The XML based web service technology allows running either simple queries or more complicated functions that are stored on the server side, where the data is.

The virtual observatory concept is adapted to the ETOMIC system. The majority of the experimental data collected in the system are inserted into the data repositories and can be reached through both the nmVO web graphical user interface and the web service interface. For the data insertion tasks a C++ and a Python language API has been developed. This API is integrated into the new APE nodes' software, so that all measurement data from the APE boxes are automatically copied into the nmVO data repository. In addition to the historical raw data collections and the evaluated results of periodic measurements also non-ETOMIC traceroute logs and topology data, one-way delay values, queuing delay tomography data, available bandwidth results, router interface clustering and IP geolocalization data can be found in the archive. The users can freely construct SQL queries to intelligently filter the raw data for their purposes. In addition to that, for some of the most common processing tasks the nmVO offers built-in server side functions, e.g. for histogram calculation or fitting the parameters of Weibull distributions.

The nmVO can be accessed through the ETOMIC website's **Database / CasJobs Query Interface** menu [1] and via web services (for which the WSDL can be requested from the server) for client applications.

All the user information, the structure of the archive, the history, queue limitation settings, etc. are stored in a separate management database. This makes it possible to search among the queries in the history.

6 Conclusion

In this paper we presented the enhanced ETOMIC network measurement infrastructure. We described the key components of the architecture and the new features of the Central Management System. The improved system kernel includes support for periodic measurements and the federated usage of the high precision ETOMIC nodes and the numerous PlanetLab nodes. Besides the kernel development novel hardware components have been developed and deployed. New lightweight measurement equipments have been installed that provide measurement services which can be remotely called by user applications via web services. The system now also includes well structured data repositories to archive and share the experimental data. Periodic measurement data can be polled with customizable pre-defined queries, while the nmVO framework gives full SQL access to its archive. The recent developments make ETOMIC an easy to use experimental facility with versatile features for network research.

Acknowledgements

The authors thank the support of the EU ICT OneLab2 Project (No.224263), the EU ICT MOMENT Project (No.215225) and the National Office for Research and Technology (NAP2005/KCKHA005).

References

1. The ETOMIC website: <http://www.etomic.org>
2. D. Morato, E. Magana, M. Izal, J. Aracil, F. Naranjo, F. Astiz, U. Alonso, I. Csabai, P. Haga, G. Simon, J. Steger, G. Vattay: *ETOMIC: A testbed for universal active and passive measurements*, IEEE TRIDENTCOM 2005, Best Testbed Award, p283-289, 23-25 Feb. 2005, Trento, Italy
3. E. Magana, D. Morato, M. Izal, J. Aracil, F. Naranjo, F. Astiz, U. Alonso, I. Csabai, P. Haga, G. Simon, J. Steger, G. Vattay: *The European Traffic Observatory Measurement Infrastructure (ETOMIC)* IEEE IPOM 2004, p165-169, 11-13 October 2004, Beijing, China (2004)
4. OneLab - Future Internet Test Beds: <http://www.onelab.eu>
5. L. Peterson, T. Anderson, D. Culler and T. Roscoe: *A Blueprint for Introducing Disruptive Technology into the Internet* Workshop on Hot Topics in Networks, Oct. 2002. <http://www.planet-lab.org>
6. P. Mátray, I. Csabai, P. Haga, J. Steger, L. Dobos, G. Vattay: *Building a Prototype for Network Measurement Virtual Observatory* ACM SIGMETRICS - MineNet 2007, 12 June 2007, San Diego, CA, USA (2007)
7. B. Augustin, X. Cuvellier, B. Orgogozo, F. Viger, T. Friedman, M. Latapy, C. Magnien and R. Teixeira: *Avoiding traceroute anomalies with Paris traceroute* Internet Measurement Conference, Oct. 2006.