African Journal of Information and Communication Technology, Vol. 2, No. 1, March 2006

An Overview of Internet Measurements: Fundamentals, Techniques, and Trends

Artur Ziviani
National Laboratory for Scientific Computing (LNCC)
Av. Getúlio Vargas, 333 – 25651-075, Petrópolis, RJ, Brazil
Email: ziviani@lncc.br

***Abstract*— The Internet presents great challenges to the characterization of its structure and behavior. Different reasons contribute to this situation, including a huge user community, a large range of applications, equipment heterogeneity, distributed administration, vast geographic coverage, and the dynamism that are typical of the current Internet. In order to deal with these challenges, several measurement-based approaches have been recently proposed to estimate and better understand the behavior, dynamics, and properties of the Internet. The set of these measurement-based techniques composes the Internet Measurements area of research. This overview paper covers the Internet Measurements area by presenting measurement-based tools and methods that directly influence other conventional areas, such as network design and planning, traffic engineering, quality of service, and network management.**

***Index Terms*— Internet Measurements, Measurement-based techniques.**

I. Introduction

ABOUT a decade ago, the Internet has begun its transformation from a structure restricted to the scientific community into a fundamental component of the new information society. Possibly, the most important consequence of the great success of the Internet is that the common goal that used to guide its players no longer holds. Users, commercial access providers, governments, telecommunication operators, and content providers have interests that can be in opposition to each other, leading to a situation where they live in a tussle [14]. For example, commercial providers need to be interconnected to obtain universal connectivity, even if they are often fierce competitors. The heterogeneity and the distributed administration of this scenario, allied to the vast geographic coverage and to the dynamism that are typical of today's Internet, impose great challenges to the characterization of the structure and behavior of the Internet as a whole [28].

Currently, the Internet has a huge community of over 800 million users with an expansion rate of 126% between 2000 and 2005 [38]. This ever-increasing user community makes use of a large range of applications. These applications generate a highly diversified traffic and require new quality services. As a function of this diversity, providers, users, and operators become aware of the need to better understand the dynamic structure and behavior of the network.

The seminal work by Paxson [61] has introduced a measurement-based approach to characterize the traffic dynamics of the Internet. Other measurement-based works have characterized the self-similar nature of network traffic in local networks [48], in wide-area networks [63], and for WWW traffic [16]. Taking into account the concepts of long-range dependence and self-similarity has significantly influenced Internet traffic modeling over the last decade [27], [44]. The work by the Faloutsos brothers [26] has also had a large impact on modeling when they suggested that the apparently random shape of the Internet topology actually followed power laws. This implies the possibility of estimating important parameters like the average number of neighbors and influences on protocol design and analysis. This feature can also be used to generate more realistic synthetic network topologies for simulation. Afterwards, several measurement-based approaches have been proposed to estimate and characterize different aspects of the Internet, making its behavior more observable [11], [84]. This better observability of network behavior helps to unveil some myths about characteristics and properties of the Internet [13], [72], [79]. These measurement-based techniques conceived to observe and infer different network characteristics compose what is called the Internet Measurements area of research.

This paper provides an overview of the Internet Measurements area, presenting tools and methods that have been recently proposed to infer and better understand the behavior, dynamics, and characteristics of the current Internet [6]. These tools and methods have direct influence on conventional areas, such as network design and planning, traffic engineering, quality of service (QoS) provisioning, and network management.

This paper is organized as follows. Section II introduces the fundamentals of performing measurements in the Internet. The Internet Measurements area is based on using measurements to estimate specific aspects, thus it is hard to discuss its challenges without considering specific problems. Therefore, Section III presents some measurement-based techniques to deal with representative network problems. This provides by no means an exhaustive list of the work developed in the Internet Measurements area, but we believe the presented techniques illustrate the potential of measurement-based methods in inferring Internet characteristics and behavior. Finally, Section IV summarizes the challenges found to better monitor and measure the behavior of the Internet. Perspectives and trends in the Internet Measurements area are also discussed.

II. Fundamentals

The basic operation of the Internet has been conceived with the explicit goal to minimize the complexity at its core and leave the control and adaptation at the edges. This design principle has allowed the Internet expansion to its current dimensions, but has also limited the capacity of monitoring the network dynamic behavior [34], [50]. Currently, the Internet is composed by a large number of interconnected networks administrated by different organizations that are often competitors. As a consequence, many domains are uncooperative with external performance measurements. There is a need to monitor the network so that we can deal with its increasing complexity, represented by a huge growth in extension, diversity, transmission speeds, and traffic volume. Figure 1 illustrates this increase showing the evolution in the last 15 years of the active entries used by BGP (Border Gateway Protocol) [70].

Network management commonly provides ways to monitor the status of individual nodes. SNMP (Simple Network Management Protocol) [10] allows a centralized network manager to request data from components of the network. This manager can also be alerted in the case some pre-defined events happen. The manager is limited to gather simple and individual measurements from each manageable equipment. Although routers are the ideal points to perform traffic measurements, they are in general not equipped for detailed monitoring. Router vendors avoid the addition of measurement capacity because of an eventual negative impact on packet forwarding performance. The NetFlow tool [12] is widely used by network operators and access providers. This tool samples flows to gather data about the traffic in the network. Despite being popular, NetFlow has some shortcomings that could be improved, such as having an adaptive sampling rate and a better capacity of sampling non-TCP flows [24]. As a consequence of the problems with the existing methods, several indirect methods are being proposed. The working group IPPM (IP Performance Metrics) [83] of the IETF (Internet Engineering Task Force) is dedicated to defining relevant metrics for evaluating the quality, performance, and reliability of network services.

A. *Characteristics of Measurement-based Methods*

Measurement-based approaches use either passive or active techniques [2]. Passive measurements refer to the process of monitoring the network traffic without injecting new traffic or affecting the existing one. This can happen in different network vantage points. Passive monitoring can provide detailed data about the network points where the measurements are carried out and about the traffic in transit in these points [43]. A high-performance passive monitoring system needs specialized equipment and currently the most adopted equipment for passive monitoring is the DAG card [18], originally developed in the Waikato University in New Zealand. To investigate how to deal with a potentially large amount of measurement data, there is a working group of the IETF called Packet Sampling (PSAMP) [3] dedicated to the definition of standards to perform packet sampling on network devices. The challenge is to define methods that are simple enough to be ubiquitously implemented without degrading significantly the packet forwarding rates of the current network devices. An example of passive network points is shown in Figure 2.

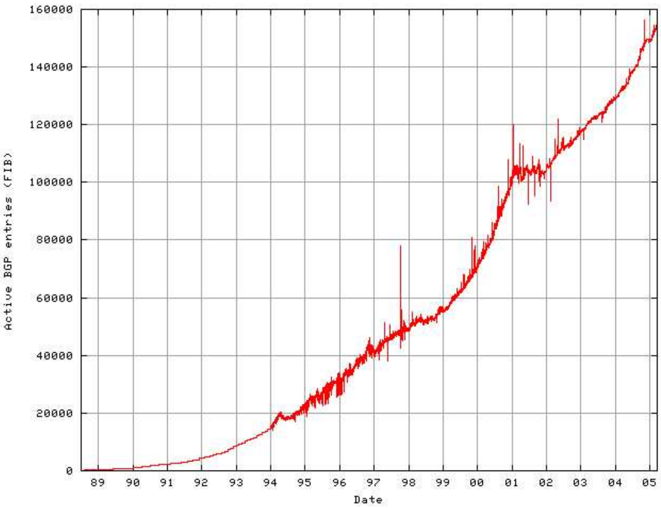

Fig. 1. Evolution of the active entries in BGP (source: [75]).

In contrast with passive measurements, active measurements send probe packets and the result of their journey through the network is monitored to estimate network characteristics. Active measurements obtain in general few information about isolated points within the network, but they can provide a representative view of the path between two points in the network. In Figure 2, active measurement probes sent from node A to node B provide information about the path between these nodes. In this example, it is assumed that the two end nodes are somehow synchronized. From the standpoint of passive measurements, only one of the passive monitoring points is able to register the passage of the probes. In performing active measurements, it is important to consider if the additional traffic introduces a bias on the obtained results or not. Hybrid scenarios may also be envisaged to better estimate a certain network characteristic [39].

Measurement methods may be classified not only in being active or passive, but they may also be differentiated by other characteristics [11]. Therefore, measurements can be:

- related to a particular packet flow or conceived to monitor the network behavior in a more general way. In the case of being related to a particular flow, measurements can be in-band, where additional fields of the packet header are used, or out-band, where additional packet probes are adopted;
- performed continuously or on demand;
- direct or indirect;
- unidirectional or bidirectional;
- composed of one or multiple data gathering points or probe measurement launching points.

The most basic and traditional tool for monitoring networks is the popular `ping`. This tool sends an `echo request` message from ICMP (Internet Control Message Protocol)

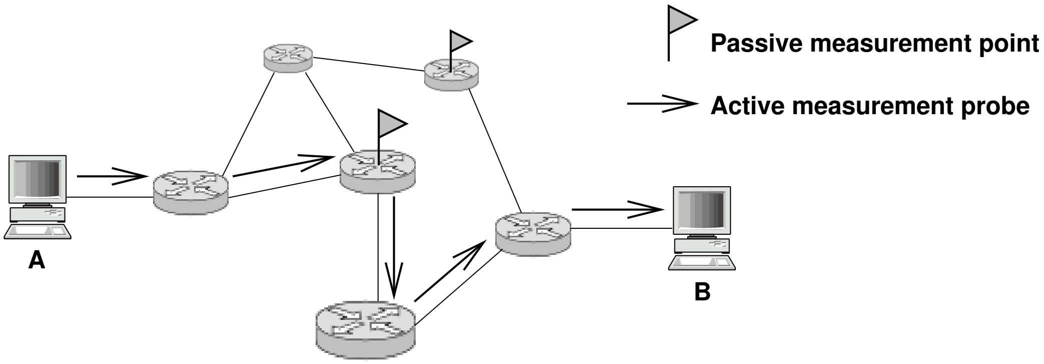

Fig. 2. Example of active and passive measurements.

to a certain network device, which in turn sends a `echo response` [68]. Besides this connectivity test, a sequence of `ping` packets also offers a simple estimation of the general network performance in the path in terms of delay and packet loss. In a single session, intervals between `ping` packets are fixed. In the case of a periodical network behavior, periodic probes may fail to accurately observe this behavior. Likewise, a periodic sampling process may synchronize with a unpredictable behavior and, as a consequence, the observed performance would be inferior than the real one. For these reasons, a Poisson sampling is recommended [60].

The `ping` tool measures the RTT (Round Trip Time), but the one-way delay is an important parameter for several applications. To measure the one-way delay, the source and destination have to be synchronized. An alternative is to synchronize the source and destination using NTP (Network Time Protocol) servers. Nevertheless, as NTP packets are distributed mixed with the remaining network traffic, synchronization errors are in the order of the observed network delays [62], leading to inaccurate measurements. Some authors [52], [85] propose methods to estimate and remove the offset and skew between clocks in nodes synchronized with NTP. The goal is to turn feasible one-way delay measurements between nodes synchronized with NTP. A straightforward solution is to adopt GPS (Global Positioning System) [23] cards for synchronization, although these need to have sight to the satellites and the cost of these cards is relatively high. Anyway, GPS cards are the current adopted solution to have precise synchronization among dispersed monitoring points in large measurement projects. Recently, in [59], authors propose an alternative software clock to enhance measurement accuracy without using GPS cards.

B. Traffic Classification and Characterization

An Internet link carries a mix of flows generated by a wide range of applications and transmitted using different transport protocols, in particular TCP and UDP [7]. The literature in Internet measurements often makes use of analogies with animals to classify network flows [4], [77]. Considering the flow size, large flows, like file transfers, are called elephants. In the other hand, small flows with low volume of data, like `http` requests, are called mice. The elephant flows may be two to three orders of magnitude larger than the mice flows [25], [56]. The fundamental difference between elephants and mice refers to the fact that a TCP session characterized as an elephant is affected by the slow start phase of the TCP congestion control mechanism. As a consequence, the behavior of an elephant flow is conditioned by the TCP congestion control. In contrast, mice flows are not controlled by this mechanism, because they are totally transmitted before the TCP is able to apply its congestion control.

As an alternative to the flow classification in terms of size (in octets), one can also classify flows in terms of their lifetime (in seconds). In [5], authors propose new criteria for flow classification adopting their lifetime as a basis. In one hand, a large amount of flows are identified as very fast, with a duration of less than 2 s. These fast flows, called dragonflies, represent at least 45% of the flows in the observed links. Close to 98% of the observed flows are less than 15 minutes long. In the other hand, the remaining 2% of flows reach durations of hours or even days. This long duration flows are called tortoises. Although the tortoise flows represent only 2% of the total number of flows, they carry 40% to 50% of the total volume of traffic. It is also shown in [5] that the size of flows in octets and the lifetime of them are independent dimensions, suggesting that size and lifetime of flows are both important to the understanding of network behavior.

III. Measurement-based Techniques

In this section we provide a brief overview of some areas where measurement-based approaches are being proposed. This is not meant to be an exhaustive list, but it is intended to show some representative measurement-based work able to illustrate the potential of these to deal with network problems.

A. Bandwidth Estimation

Network administrators that have privileged access to a router or switch connected to a link of interest may directly measure some parameters related to the bandwidth on that link. This can be done by SNMP. Nevertheless, this access is typically available only to administrators and not to end users. The end user can only *estimate* the link bandwidth using network measurements. Even network administrators, with privileged access to some routers, may need to determine the bandwidth between routers under their control and external

TABLE I

SOME BANDWIDTH ESTIMATION TOOLS.

Tool	Metric	Method	Reference
`pathchar`	per-link capacity	VPS	[40]
`clink`	per-link capacity	VPS	[21]
`bprobe`	end-to-end capacity	PPTD	[9]
`nettimer`	end-to-end capacity	PPTD	[45]
`pathrate`	end-to-end capacity	PPTD	[20]
`sprobe`	end-to-end capacity	PPTD	[71]
`pathload`	available bandwidth	SLoPS	[41]
`IGI/PTR`	available bandwidth	SLoPS	[36]

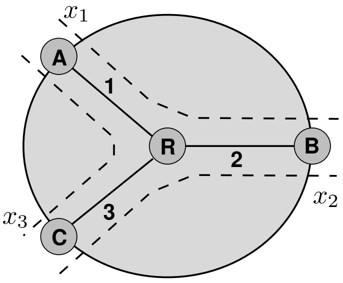

Fig. 3. Example of traffic matrix estimation.

ones. In this case, these administrators also use estimation techniques that are based on network measurements.

In [69], authors define metrics associated with the bandwidth estimation. First, it is established a difference between individual link bandwidth and the bandwidth of a sequence of links, *i.e.* the end-to-end path. Second, the metrics are capacity and available bandwidth. Capacity refers to the maximum bandwidth that can be reached in a link or path. Available bandwidth is the maximum idle bandwidth in a link or path. The identification of the lowest capacity along a path, *i.e.* the bottleneck link, also draws attention from different researchers [35].

There are three main techniques to estimate bandwidth: variable packet size (VPS) probing, packet pair/train dispersion (PPTD), and self-loading of periodic streams (SLoPS). The first technique infers the capacity of individual links. The second one estimates the end-to-end capacity. The third technique infers the end-to-end available bandwidth. In general, these techniques assume that, during the measurement process, the end-to-end path remains the same and the traffic is stationary. Dynamic changes in routing and load may raise problems in any of these methods. In [42], several issues on estimating bandwidth are pointed out. Table I presents a list of some bandwidth estimation tools. A recent analysis of public available tools for estimating bandwidth can be found in [74].

B. *Traffic Matrix Estimation*

It is often hard to directly measure the traffic matrix of a large operational IP network because of the costly additional infrastructure needed [57]. Therefore, in general, complete traffic matrices are not available to large network operators. Nevertheless, measurements about the total load at each individual link are readily available in IP networks by using regular management tools. Thus, to estimate the traffic matrix in the IP network of a large operator one needs to estimate the end-to-end traffic demands within the domain from the individual link loads. This problem of estimating a traffic matrix from partial information about individual link loads is commonly called network tomography and has received a lot of attention from the measurement research community.

The problem of traffic matrix estimation can be formalized in the following way [51]. Let c denote the number of origin-destination (OD) pairs within a network domain. If this domain has n nodes of interest at its borders, then $c = n(n-1)$. The OD pairs are ordered in a vector $\mathbf{x}$, where $x_j \in \mathbf{x}$ is the traffic volume transmitted in the OD pair j. Let $\mathbf{y} = [y_1, \ldots, y_r]^T$ be the vector that represents the traffic volume at each link individually. The element y_l indicates the traffic volume at link l and r denotes the number of links in the domain. Vectors $\mathbf{x}$ and $\mathbf{y}$ are related through a routing matrix $\mathbf{A}_{r \times c}$. Matrix $\mathbf{A}$ is composed by $\{0, 1\}$ with rows representing the network links and columns representing the OD pairs. The element $a_{ij} = 1$ indicates that link i belongs to the path associated with OD pair j, and $a_{ij} = 0$ otherwise. Therefore, the OD pairs are related to the individual traffic volumes in accordance with the following linear relation:

$$\mathbf{y} = \mathbf{A}.\mathbf{x} \tag{1}$$

To better understand the composition of $\mathbf{y}$, $\mathbf{A}$, and $\mathbf{x}$, observe Figure 3 that illustrates the problem of estimating traffic matrices. In this figure there are three nodes of interest A, B, and C interconnected using router R through links 1, 2, and 3. The information about the individual load in these links is available and compose the vector $\mathbf{y} = [y_1, y_2, y_3]^T$. The OD pairs that are the elements of the traffic matrix $\mathbf{x}$ are represented by the slashed lines in Figure 3[1]. The problem is to estimate the traffic matrix $\mathbf{x}$ whose elements are x_1 that represents the OD pair between nodes A and B, x_2 that indicates the OD pair between nodes B and C, and x_3 that denotes the OD pair between nodes A and C. For instance, $y_1 = x_1 + x_3$. Thus, the relation $\mathbf{y} = \mathbf{A}.\mathbf{x}$ in the case illustrated in Figure 3 is given by:

$$\begin{bmatrix} y_1 \\ y_2 \\ y_3 \end{bmatrix} = \begin{bmatrix} 1 & 0 & 1 \\ 1 & 1 & 0 \\ 0 & 1 & 1 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} \tag{2}$$

The routing matrix in IP networks can be obtained by gathering data from the underlying routing protocols and the computation of the shortest paths between all OD pairs. The traffic volumes at individual links are available through the use of SNMP. Hence, the problem consists of computing $\mathbf{x}$, *i.e.* the set of OD pairs that reproduce the traffic volume at the links in the closest possible manner. The problem associated with Equation (1) is highly undetermined because

[1] For the sake of simplicity, it is considered in this example the load of an OD pair as the bidirectional total load, *i.e.* x_1 includes the traffic from A to B, and vice-versa. In a real traffic matrix, we need to consider unidirectional OD pairs given that routing is in general asymmetric.

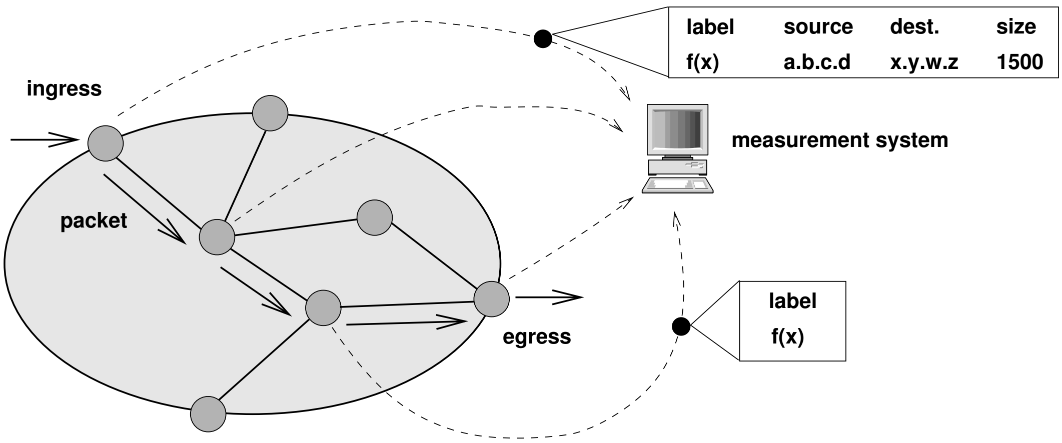

Fig. 4. Trajectory sampling.

in practice the number of OD pairs is often much larger than the number of links, $r \ll c$. This means that there are infinite possible solutions to finding $\mathbf{x}$. Medina *et al.* [51] compare different methods to estimate traffic matrices. The effects of routing changes on the variation of estimated traffic matrices is investigated in [82].

C. Traffic Sampling and Anomaly Diagnosis

Network anomalies are defined as significant and unlikely changes in traffic patterns at one or multiple links [1]. The diagnosis of these anomalies involves detecting, identifying, and quantifying them. Independent of being intentional or not, anomalies are worth analyzing because of two reasons. First, anomalies may cause network congestion and quickly consume router resources, thus making their identification crucial from the viewpoint of network operators. Second, some anomalies not necessarily affect network performance, but they may have important impact on clients or end users. The anomaly diagnosis presents great challenges because one must extract anomalous patterns from large volumes of data and the anomaly causes can be highly diversified. Examples of anomalies include distributed denial of service attacks, misconfiguration of routers, or results from the modification of policies in BGP routing.

A method for detecting, identifying, and quantifying anomalies is proposed in [47]. Detection consists of determining the points in time in which the network is affected by an anomaly. Identification involves the classification of the detected anomaly out of a set of known anomaly patterns. Quantification measures the importance of the anomaly by estimating how much anomalous traffic is in the network.

In detecting an anomaly, an interesting functionality would be the capacity of tracking the trajectory of packets that compose an anomalous traffic within a domain. This measurement-based capacity makes the network more resilient to failures and to the presence of anomalies. In [22], authors propose a method to sample packet trajectories in a network domain. The sampling methodology selects a subset of packets, but if a packet is selected at a link, it is selected at all links the packet traverses. Through the network, each packet indicates implicitly if it should be sampled or not because of its invariant part. A hash function is applied in each router at this invariant part of packets. Only the packets whose hash result falls into a certain interval are selected for sampling. In this way, if the same hash function is adopted throughout the domain to select packets for sampling, then there is a guarantee that either the packet is selected at all links it traverses or the packet is never selected. Therefore, this method enables the collection of trajectory samples of a subset of packets.

Sampled packets also generate a label using a second hash function to identify each sampled packet. Assuring the uniqueness of labels at least for a minimum time period allows the observation of the subset of links that has been traversed by a particular packet because these links would have reported the passage of the same identification label. Figure 4 presents an example of trajectory sampling. Solid arrows represent the path through the domain taken by a packet whose invariant parts trigger the sampling process. Using the second hash function to identify the packet, the routers send the resulting packet label to the centralized measurement system, as indicated by the slashed line. Although this suffices to identify packet trajectories as sampled within the domain, some additional information may be useful for different measurement purposes. This additional information may include the source and destination addresses of the packet, and its size as well. Nevertheless, this information may be gathered just once per sampled packet. Hence, the ingress nodes may be configured to collect this additional information and not just the packet labels as do the remaining nodes, as illustrated in Figure 4. Note that multicast packets require no further treatment. In this case, the trajectory associated with a multicast packet is simply a tree instead of a path. A similar strategy with respect to the sampling of packet trajectories is adopted in [76] to traceback undergoing attacks.

D. Network Proximity

There is an increasing need of a means to estimate distances between nodes in the Internet [37]. In this context, distance refers to some network performance metric such as delay or bandwidth. Delay as a distance metric between nodes in the

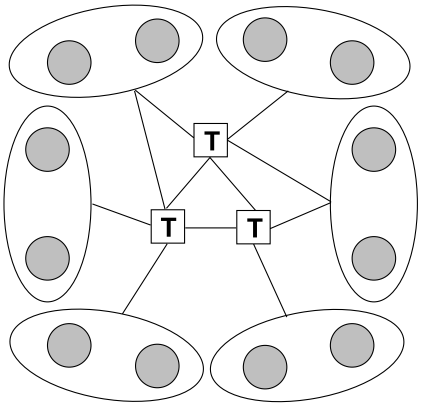

Fig. 5. IDMaps architecture.

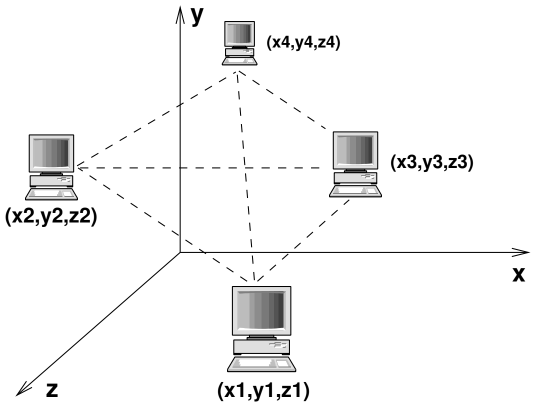

Fig. 6. Modeling the Internet as an abstract geometric space.

network provides useful information for several application and services, such as distributed website hosting, finding of the closest server, application-layer multicast, content distribution networks, and peer-to-peer (P2P) file-sharing systems.

Although nodes can measure path characteristics with tools like `ping` or `traceroute`, carrying out measurements before each interaction on the Internet would certainly lead to an overload on both end nodes and the network itself. Therefore, the goal is to evaluate network proximity in terms of delay in a scalable way, but without the need to perform direct measurements between nodes. Measurements of the delay distance between nodes usually consist of the minimum observed delay to avoid taken into account buffering delays that probe packets may encounter at intermediate routers.

1) Delay estimation between two network points: IDMaps [30] was the first proposal of a global architecture to estimate network distances between nodes in the Internet. IDMaps defines some systems, called tracers, that are distributed in the Internet with the goal of having tracers relatively close to any address prefix (Figure 5). The distances between the tracers are measured as well as the distance between the networks represented by the address prefixes and the closest tracer. The distance between any two address prefixes is estimated as the sum of the distances from each address prefix to the closest tracer and the distance between the tracers.

The quality of the resulting estimated distance depends on the number of adopted tracers and on their localization. Therefore, there is a trade-off between improving the quality of results at the expense of more measurements. The IDMaps approach results in a set of distances in the order of $B^2 + P$, where B is the number of tracers and P is the number of address prefixes. The number of address prefixes is in the order of 150,000 as of March 2005 [75]. Thus, if the number of tracers B is limited to a few hundreds, the total volume of distances to be managed becomes feasible. The system operates in a client-server architecture where HOPS (HOst Proximity Service) servers provide the distance between two arbitrary nodes using measurements done by the IDMaps architecture. In the evaluation of IDMaps presented in [30], it is shown that the number of tracers needed to obtain satisfactory results is relatively small, since using only 0.2% of nodes as tracers provides a correct answer in 90% of the observed cases.

2) Approaches based on coordinate systems: As an alternative to the client-server architecture of IDMaps, other proposals have emerged to estimate network proximity based on a P2P model. This model has a larger potential for scalability when compared to the client-server model. Performance bottlenecks are avoided by the absence of remote servers. Moreover, this model is consistent with P2P applications, such as file-sharing, content distribution networks, and application-layer multicast services that can significantly take benefit from information about network proximity between nodes.

GNP (Global Network Positioning) [54] was the first proposal based on P2P to estimate the network distance between two nodes in the Internet. The basic idea of GNP is to keep coordinates associated with the participating nodes in order to represent relative positions in the Internet. The network distance could then be estimated by computing a distance function with the coordinates of nodes.

In the first step, GNP adopts a small set of distributed landmark nodes to provide the reference coordinates in the resulting abstract space for other nodes. These landmarks periodically measure the distance to each other to correct their coordinates if needed. The delay distances can be measured for instance as the minimum RTT of several measurements using `ping` to avoid taking into account buffering delay at intermediate routers. The landmarks then transform the distances to each other in coordinates in the abstract space, as illustrated in Figure 6 for a hypothetical 3D space. To do so, one can adopt a method called *Multidimensional Scaling* (MDS) [80]. In a second step, common nodes can participate in the GNP system. Using the coordinates of the landmarks in the abstract space, each common node can determine its own coordinates by measuring its network distance toward the set of landmarks. In this step, the landmarks play a passive role in the process and only answer the ICMP messages from the common node that wants to join the system. This procedure is illustrated in Figure 7.

ICS (Internet Coordinate System) [49] and Virtual Landmarks [81] are two similar proposals to improve the performance in accuracy of GNP when embedding measured

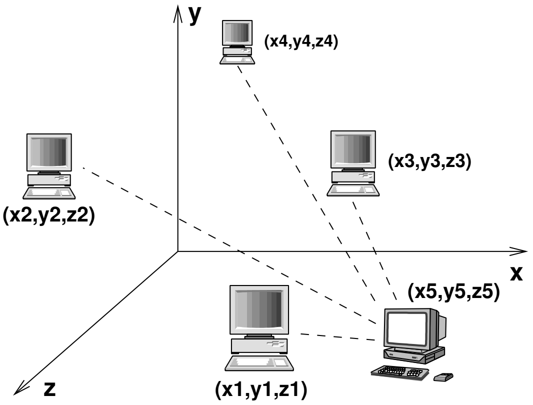

Fig. 7. Establishing the coordinates of a common node.

distances into a lower dimensional space. These proposals use the PCA (Principal Components Analysis) [8] technique to achieve this improvement. Recently, several other proposals have been inspired on the basic ideas introduced by GNP, such as the use of a P2P model and a coordinate system to estimate network proximity. These proposals include King [33], Lighthouses [67], Big-Bang Simulation (BBS) [73], Practical Internet Coordinates (PIC) [15], and Vivaldi [17].

E. Geolocation of Internet Hosts

Novel location-aware applications could be enabled by an efficient means of inferring the geographic location of Internet hosts. Examples of these location-aware applications include:

- Targeted advertising on web pages – Online consumers may have different regional preferences based on where they live. Being able to locally tailor products, marketing strategies, and contents confers a business advantage;
- Restricted content delivery – Following regional policies, a geographic location service can determine which client has access to content. Similarly, enforcement of localized regulation is enabled;
- Location-based security check – If authorized locations are known, an e-commerce transaction that is requested from elsewhere might generate warnings on atypical or unauthorized behavior of a customer.

The inference of the geographic location of Internet nodes from IP addresses constitutes a challenging problem, because there is no direct relation between the IP address of a node and its geographic location [31], [46]. This section describes techniques to estimate the geographic location of Internet hosts from their IP addresses based on delay measurements. It should be noted that in contrast with the proposals to estimate network proximity discussed in Section III-D where distances are measured in terms of delay, the techniques of this section infer the geographic location of an Internet node, thus the distances refer to the physical distances between nodes. In this section, two approaches for geolocating Internet host are presented: the discrete one and the continuous one. The former is based on finding nearby landmarks in the network infrastructure [55], [87] and the latter relies on direct multilateration to geolocate hosts [32]. The state of the art indicates that geolocation of an Internet host to the level of a metropolitan area is often feasible.

1) Discrete Geolocation: In a discrete system, the geographic location of nodes is inferred by comparing the delay pattern of each landmark and the one observed for the target host. The landmark that presents the most similar delay pattern with respect to the one of the target host provides the location estimation of that host. This is the base of the GeoPing technique [55], which adopts an empirical approach based on the observation that hosts sharing similar delays to other fixed hosts tend to be near each other geographically.

We formalize the problem of inferring a host location from delay measurements as follows. Consider a set $\mathcal{L} = \{L_1, L_2, \ldots, L_K\}$ of K landmarks. Landmarks are reference hosts with a well-known geographic location. Consider a set $\mathcal{P} = \{P_1, P_2, \ldots, P_N\}$ of N probe machines. Figure 8 illustrates the steps in inferring a host location from delay measurements, which are detailed along this section. Dotted lines represent the measurements taken by the probe machines while the solid lines indicate information exchange. The probe machines periodically determine the network delay, which is actually the minimum delay of several measurements, to each landmark (Figure 8(a)). Therefore, each probe machine P_x, $1 \leq x \leq N$, keeps a delay vector $\mathbf{d}_x = [d_{1x}, d_{2x}, \ldots, d_{Kx}]^T$, where d_{ix} is the delay between the probe machine P_x and the landmark $L_i \in \mathcal{L}$. Suppose one wants to determine the location of a given target host τ. A location server that knows the landmark set $\mathcal{L}$ and the probe machine set $\mathcal{P}$ is then contacted. The location server asks the N probe machines to measure the delay to host τ (Figure 8(b)). Each probe machine P_x, $1 \leq x \leq N$, returns a delay vector $\mathbf{d}'_x = [d_{1x}, d_{2x}, \ldots, d_{Kx}, d_{\tau x}]^T$, i.e., the delay vector $\mathbf{d}_x$ plus the just measured delay to host τ (Figure 8(c)). After receiving the delay vectors from the N probe machines, the location server is able to construct the delay matrix $\mathbf{D}_{(K+1)\times N}$:

$$\mathbf{D} = \begin{bmatrix} d_{11} & d_{12} & \ldots & d_{1N} \\ d_{21} & d_{22} & \ldots & d_{2N} \\ \vdots & \vdots & \ddots & \vdots \\ d_{K1} & d_{K2} & \ldots & d_{KN} \\ d_{\tau 1} & d_{\tau 2} & \ldots & d_{\tau N} \end{bmatrix} \qquad (3)$$

The delay vectors gathered by the demanding location server from the probe machines correspond to the columns of the delay matrix $\mathbf{D}$. The location server then compares the lines of the delay matrix $\mathbf{D}$ to estimate the location of host τ. To infer the location of host τ, the landmark L presenting the most similar delay pattern with respect to the delay pattern of host τ is determined. The corresponding location of the landmark L is the location estimation of host τ (Figure 8(d)). The delay matrix $\mathbf{D}$ combined with the knowledge of the location of the landmarks of the set $\mathcal{L}$ compose a delay map recording the relationship between network delay and geographic location. Practical results of measurement to geographically locate Internet nodes using the NIMI (National Internet Measurement Infrastructure) [64] platform are presented in [86].

In [87], some techniques are proposed to improve the geolocation of Internet hosts using the discrete system.

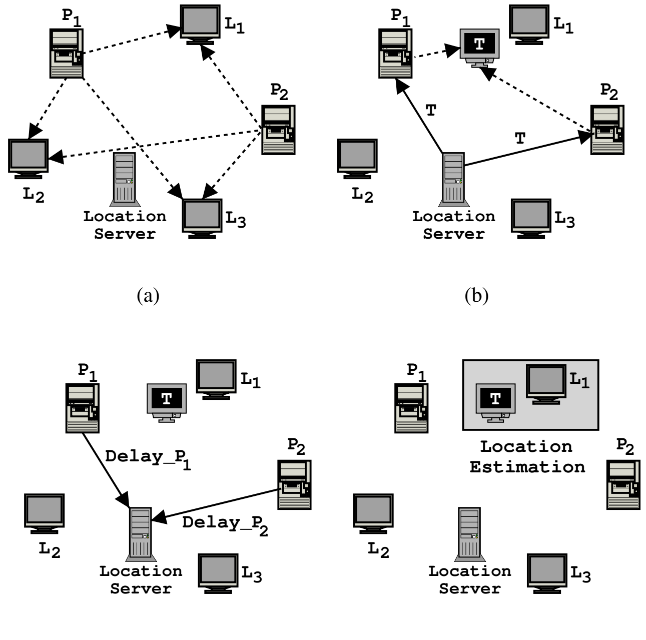

(a) (b) (c) (d)

Fig. 8. Inferring a host location from delay measurements.

Authors first investigate the correlation observed in the network between geographic distance and network delay. This correlation is weak to moderate if considered globally, but it is observed that this correlation becomes stronger in regions with rich connectivity. The term rich, or poor, applied to connectivity represents the diversity of available connectivity and the transit options found in certain regions at either router or autonomous system levels. From an environment with rich connectivity, it is expected more routing options that can roughly approximate the direct geographic path between source and destination. Two key points that influence the accuracy of the discrete system are identified: the placement of landmarks and probe machines, and how efficient is the evaluation of similarity between the delay patterns. Hence, in [87], it is suggested to improve the performance of the discrete system to estimate the geolocation of Internet hosts in two ways: (i) to strategically place landmarks and probe machines; and (ii) to select models to better evaluate the similarity between the delay patterns of the landmarks and the one observed for the target host.

2) Continuous Geolocation: Previous works [55], [87] use the position of landmarks, which have a well-known geographic location, as possible location estimates of the target hosts. This leads to a discrete space of answers, *i.e.* the number of answers is equivalent to the number of landmarks. This may lead to inaccurate results since the closest landmark may still be far from the target.

CBG (Constraint-Based Geolocation) [32] is proposed to overcome the limitation of a discrete space system. This is achieved by using multilateration, which refers to the process of estimating a position using a sufficient number of distances to some fixed points. As a result, multilateration establishes a continuous space of answers instead of a discrete one. CBG adopts a set of landmarks to estimate the location of other Internet nodes. The fundamental idea is that given the geographic distances to a certain target node from the landmarks, a location estimation is feasible using multilateration, as done in GPS.

A key element of CBG is its ability to accurately transform delay measurements into distance constraints. The starting point is the fact that digital information travels along fiber optic cables at almost exactly 2/3 the speed of light in a vacuum [65]. This means that any particular delay measurement immediately provides an *upper bound* on the great-circle distance between the endpoints. The upper bound is the delay measurement divided by the speed of light in fiber. Looking at this from the standpoint of a particular pair of endpoints, we can reason that there is some theoretical minimum delay for packet transmission that is dictated by the great-circle distance between them. Therefore, no matter the reason (*e.g.* queuing delays, violations of the triangle inequality, absence of great-circle paths between hosts, and so on), the actual measured delay between them involves only an *additive* distortion.

However, if CBG were to use simple delay measurements directly to infer distance constraints, it would not be very accurate. For accurate results, it is important to estimate and remove as much of the additive distortion as possible. CBG does this by self-calibrating the delay measurements taken from each measurement point in a distributed manner. After self-calibration, CBG can more accurately transform a set of measured delays to a target into distance constraints. CBG then uses multilateration with these distance constraints to establish a geographic region that contains the target host. Given this target region, a reasonable "guess" as to the host's location is at the region's centroid, which is what CBG uses as a point estimate of the target's position. It should be noted that, in contrast with other proposals, CBG associates a confidence region to each location estimate. This allows location-aware applications to decide if the provided location estimate has sufficient resolution with respect to their particular needs.

Figure 9 illustrates the multilateration in CBG using the set of landmarks $\mathcal{L} = \{L_1, L_2, L_3\}$ in the presence of some additive distance distortion due to imperfect measurements. Each landmark L_i intends to infer its geographic distance constraint to a target host τ with unknown geographic location. Nevertheless, the inferred geographic distance constraint is actually given by $\hat{g}_{i\tau} = g_{i\tau} + \gamma_{i\tau}$, *i.e.* the real geographic distance $g_{i\tau}$ plus an additive geographic distance distortion represented by $\gamma_{i\tau}$. This purely additive distance distortion $\gamma_{i\tau}$ results from the eventual presence of some additive delay distortion. As a consequence of having additive distance distortion, the location estimation of the target host τ should lie somewhere within the gray area (*cf.* Figure 9) that corresponds to the intersection of the overestimated geographic distance constraints from the landmarks to the target host.

IV. Summary and Outlook

Internet protocols have not been originally conceived to support detailed performance measurements. For this reason,

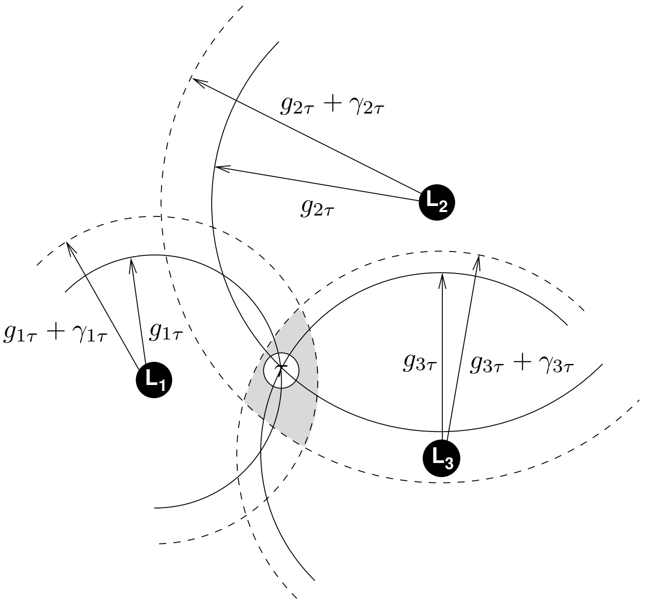
$g_{2\tau} + \gamma_{2\tau}$
$g_{2\tau}$
L_2
$g_{1\tau} + \gamma_{1\tau}$
$g_{1\tau}$
L_1
$g_{3\tau}$
$g_{3\tau} + \gamma_{3\tau}$
L_3

Fig. 9. Multilateration with geographic distance constraints.

developers and researchers need to investigate means to indirectly measure network characteristics and behavior. This review paper introduces the fundamentals and describes some methods in the area of Internet Measurements. These measurement-based methods deal with problems such as bandwidth estimation, traffic matrix estimation, traffic sampling and anomaly diagnosis, network proximity evaluation, and geolocation of Internet hosts. Other areas also receive large attention from active research efforts such as topology inference in the router level [78] or autonomous system level [19], traffic forecasting [58], network support for grid computing [88], and identification and characterization of applications [53], among others.

Despite of the advancements in the Internet Measurements area in recent years, the gathering, sampling, interpretation, and modeling of empirical Internet data still pose challenging problems. The first challenge is that several aspects of the Internet are ever-changing. For example, `http` traffic has grown from zero in 1995 to more than 80% of the total traffic in the majority of links in 2000. Currently, the proportion of `http` traffic seems to be decreasing in the majority of links while there is an increasing presence of P2P traffic [29].

The global scale of the Internet also imposes great challenges to measurement projects, because quite often the composition of traffic and its behavior are dependent on location and characteristics of particular groups of users. As a consequence, observed results in a single location may not be representative of the Internet as a whole. Therefore, measurements need to be performed from multiple points to obtain a more representative view of the big picture. Platforms like NIMI (National Internet Measurement Infrastructure) [64] and PlanetLab [66] provide distributed infrastructure that can be applied to carry out large-scale measurement experiments.

Making the Internet more observable may be the first step in the direction of having a more efficient monitoring of the network. Just collecting huge volumes of measurement data is not efficient without the development of advanced tools to process such a volume of data and provide a basis to the design of more efficient applications and services. New measurement-based techniques as well as new methods for sampling and inferring characteristics of the Internet are in need and may open promising perspectives to novel research activity based on Internet Measurements.

ACKNOWLEDGMENT

The author is grateful to many people for different reasons that are expressed in the following. Antônio Tadeu A. Gomes (LNCC) has provided helpful comments on preliminary versions of this text. Timur Friedman (LIP6/CNRS), Andrew Adams (PSC), and Vern Paxon (ICIR) have enabled the use of the NIMI platform by the author. Otto Carlos M. B. Duarte (COPPE/UFRJ), José Ferreira de Rezende (COPPE/UFRJ), Bruno Schulze (LNCC), Bamba Gueye (LIP6/CNRS), Serge Fdida (LIP6/CNRS), and Mark Crovella (Boston University) are co-authors of different papers in the area of Internet Measurements and the lessons taken from these works have been very useful for writing this overview on this topic.

REFERENCES

[1] BARFORD, P., KLINE, J., PLONKA, D., AND RON, A. A signal analysis of network traffic anomalies. In *Proc. of ACM/SIGCOMM Internet Measurement Workshop – IMW 2002* (Marseille, France, Nov. 2002).

[2] BARFORD, P., AND SOMMERS, J. Comparing probe-based and router-based packet-loss measurement. *IEEE Internet Computing 8*, 5 (Sept. 2004), 50–56.

[3] BIERMAN, A., AND QUITTEK, J. *Packet Sampling (PSAMP)*, 2001. `http://www.ietf.org/html.charters/psamp-charter.html`.

[4] BROWNLEE, N. Some observations of Internet stream lifetimes. In *Proc. of the Passive and Active Measurement Workshop - PAM'2005* (Boston, MA, USA, Apr. 2005).

[5] BROWNLEE, N., AND CLAFFY, K. C. Understanding Internet traffic streams: dragonflies and tortoises. *IEEE Communications Magazine 40*, 10 (Oct. 2002), 110–117.

[6] BROWNLEE, N., AND CLAFFY, K. C. Internet measurement. *IEEE Internet Computing 8*, 5 (Sept. 2004), 30–33.

[7] BROWNLEE, N., MILLS, C., AND RUTH, G. Traffic flow measurement: Architecture. *RFC 2722* (Oct. 1999).

[8] BRYANT, F. B., AND YARNOLD, P. R. *Reading and Understanding Multivariate Statistics*. APA Press, 1998, ch. Principal-Components Analysis and Exploratory and Confirmatory Factor Analysis, pp. 99–136.

[9] CARTER, R. L., AND CROVELLA, M. Measuring bottleneck link speed in packet-switched networks. *Performance Evaluation 27-28* (Oct. 1996), 297–318.

[10] CASE, J. D., FEDOR, M., SCHOFFSTALL, M. L., AND DAVIN, J. R. Simple network management protocol (SNMP). *RFC 1157* (May 1990).

[11] CHEN, T. M. Increasing the observability of Internet behavior. *Communications of the ACM 44*, 1 (Jan. 2001), 93–98.

[12] CISCO. *NetFlow*, 1999. `http://www.cisco.com/warp/public/732/Tech/nmp/netflow/`.

[13] CLAFFY, K. C. Internet measurement: myths about Internet data. Talk at NANOG24 Meeting, Feb. 2002. `http://www.caida.org/outreach/presentations/Myths2002/`.

[14] CLARK, D. D., WROCLAWSKI, J., SOLLINS, K. R., AND BRADEN, R. Tussle in cyberspace: Defining tomorrow's Internet. In *Proc. of the ACM SIGCOMM'2002* (Pittsburgh, PA, USA, Aug. 2002).

[15] COSTA, M., CASTRO, M., ROWSTRON, A., AND KEY, P. PIC: Practical Internet coordinates for distance estimation. In *Proc. of the IEEE International Conference on Distributed Computing Systems – IEEE ICDCS'2004* (Tokyo, Japan, Mar. 2004).

[16] CROVELLA, M. E., AND BESTAVROS, A. Self-similarity in world wide web traffic: evidence and possible causes. *IEEE/ACM Transactions on Networking 5*, 6 (Dec. 1997), 835–846.

[17] Dabek, F., Cox, R., Kaashoek, F., and Morris, R. Vivaldi: A decentralized network coordinate system. In *Proc. of the ACM SIGCOMM'2004* (Portland, OR, USA, Aug. 2004).

[18] DAG. DAG cards – Endace measurement systems, Apr. 2001. http://www.endace.com.

[19] Dimitropoulos, X., Krioukov, D., and Riley, G. Revisiting Internet AS-level topology discovery. In *Proc. of the Passive and Active Measurement Workshop - PAM'2005* (Boston, MA, USA, Apr. 2005).

[20] Dovrolis, C., Ramanathan, P., and Moore, D. Packet dispersion techniques and a capacity estimation methodology. *IEEE/ACM Transactions on Networking 12*, 6 (Dec. 2004), 963–977.

[21] Downey, A. B. Using pathchar to estimate Internet link characteristics. In *Proc. of the ACM SIGCOMM'99* (Cambridge, MA, USA, Sept. 1999).

[22] Duffield, N., and Grossglauser, M. Trajectory sampling for direct traffic observation. *IEEE/ACM Transactions on Networking 9*, 3 (June 2001), 280–292.

[23] Enge, P., and Misra, P. Special issue on global positioning system. *Proceedings of the IEEE 87*, 1 (Jan. 1999), 3–15.

[24] Estan, C., Keys, K., Moore, D., and Varghese, G. Building a better NetFlow. In *Proc. of the ACM SIGCOMM'2004* (Portland, OR, USA, Aug. 2004).

[25] Estan, C., and Varghese, G. New directions in traffic measurement and accounting. In *Proc. of the ACM SIGCOMM'2002* (Pittsburgh, PA, USA, Aug. 2002).

[26] Faloutsos, M., Faloutsos, P., and Faloutsos, C. On power-law relationships of the Internet topology. In *Proc. of the ACM SIGCOMM'99* (Cambridge, MA, USA, Sept. 1999).

[27] Figueiredo, D. R., Liu, B., Feldmann, A., Misra, V., Towsley, D., and Willinger, W. On TCP and self-similar traffic. *Performance Evaluation* (2005). Special issue on Long Range Dependence and Heavy Tail Distributions. To appear.

[28] Floyd, S., and Paxson, V. Difficulties in simulating the Internet. *IEEE/ACM Transactions on Networking 9*, 4 (Aug. 2001), 392–403.

[29] Fraleigh, C., Moon, S., Lyles, B., Cotton, C., Khan, M., Moll, D., Rockell, R., Seely, T., and Diot, C. Packet-level traffic measurements from the Sprint IP backbone. *IEEE Network 17*, 6 (Nov. 2003), 6–16.

[30] Francis, P., Jamin, S., Jin, C., Jin, Y., Raz, D., Shavitt, Y., and Zhang, L. IDMaps: A global Internet host distance estimation service. *IEEE/ACM Transactions on Networking 9*, 5 (Oct. 2001), 525–540.

[31] Freedman, M. J., Vutukuru, M., Feamster, N., and Balakrishnan, H. Geographic locality of IP prefixes. In *Proc. of ACM/SIGCOMM Internet Measurement Conference – IMC 2005* (Berkeley, CA, USA, Oct. 2005).

[32] Gueye, B., Ziviani, A., Crovella, M., and Fdida, S. Constraint-based geolocation of Internet hosts. In *Proc. of ACM/SIGCOMM Internet Measurement Conference – IMC 2004* (Taormina, Italy, Oct. 2004).

[33] Gummadi, K. P., Saroiu, S., and Gribble, S. D. King: Estimating latency between arbitrary Internet end hosts. In *ACM Internet Measurement Workshop 2002* (Marseille, France, Nov. 2002).

[34] Habib, A., Khan, M., and Bhargava, B. Edge-to-edge measurement-based distributed network monitoring. *Computer Networks 44*, 2 (Feb. 2004), 211–233.

[35] Hu, N., Li, L., Mao, Z. M., Steenkiste, P., and Wang, J. A measurement study of Internet bottleneck. In *Proc. of the IEEE INFOCOM'2005* (Miami, FL, USA, Mar. 2005).

[36] Hu, N., and Steenkiste, P. Evaluation and characterization of available bandwidth probing techniques. *IEEE Journal on Selected Areas in Communications 21*, 6 (Aug. 2003), 879–894.

[37] Huffaker, B., Fomenkov, M., Plummer, D. J., Moore, D., and k claffy. Distance metrics in the Internet. In *Proc. of the IEEE International Telecommunications Symposium - ITS'2002* (Natal, Brazil, Sept. 2002).

[38] Internet World Stats. Internet usage and population statistics, Feb. 2005. http://www.internetworldstats.com/stats.htm.

[39] Ishibashi, K., Kanazawa, T., Aida, M., and Ishii, H. Active/passive combination-type performance measurement method using change-of-measure framework. *Computer Communications 27*, 9 (June 2004), 868–879.

[40] Jacobson, V. Pathchar: A tool to infer characteristics of Internet paths, Apr. 1997. http://www.caida.org/tools/utilities/others/pathchar/.

[41] Jain, M., and Dovrolis, C. End-to-end available bandwidth: Measurement methodology, dynamics, and relation with TCP throughput. In *Proc. of the ACM SIGCOMM'2002* (Pittsburgh, PA, USA, Aug. 2002).

[42] Jain, M., and Dovrolis, C. Ten fallacies and pittfalls on end-to-end available bandwidth estimation. In *Proc. of ACM/SIGCOMM Internet Measurement Conference – IMC 2004* (Taormina, Italy, Oct. 2004).

[43] Jaiswal, S., Iannaccone, G., Diot, C., Kurose, J., and Towsley, D. Inferring TCP connection characteristics through passive measurements. In *Proc. of the IEEE INFOCOM'2004* (Hong Kong, Mar. 2004).

[44] Karagiannis, T., Molle, M., and Faloutsos, M. Long-range dependence: Ten years of Internet traffic modeling. *IEEE Internet Computing 8*, 5 (Sept. 2004), 57–64.

[45] Lai, K., and Baker, M. Measuring link bandwidths using a deterministic model of packet delay. In *Proc. of the ACM SIGCOMM'2000* (Stockholm, Sweden, Aug. 2000).

[46] Lakhina, A., Byers, J., Crovella, M., and Matta, I. On the geographic location of Internet resources. *IEEE Journal on Selected Areas in Communications 21*, 6 (Aug. 2003), 934–948.

[47] Lakhina, A., Crovella, M., and Diot, C. Diagnosing network-wide traffic anomalies. In *Proc. of the ACM SIGCOMM'2004* (Portland, OR, USA, Aug. 2004).

[48] Leland, W., Taqqu, M., Willinger, W., and Wilson, D. On the self-similar nature of ethernet traffic. *IEEE/ACM Transactions on Networking 2*, 1 (Feb. 1994), 1–15.

[49] Lim, H., Hou, J. C., and Choi, C.-H. Constructing Internet coordinate system based on delay measurement. In *ACM Internet Measurement Conference 2003* (Miami, FL, USA, Oct. 2003).

[50] Mao, G. A real-time loss perfomance monitoring scheme. *Computer Communications 28*, 2 (Feb. 2005), 150–161.

[51] Medina, A., Taft, N., Salamatian, K., Bhattacharyya, S., and Diot, C. Traffic matrix estimation: Existing techniques and new directions. In *Proc. of the ACM SIGCOMM'2002* (Pittsburgh, PA, USA, Aug. 2002).

[52] Moon, S., Skelly, P., and Towsley, D. Estimation and removal of clock skew from network delay measurements. In *Proc. of the IEEE INFOCOM'99* (New York, NY, USA, Mar. 1999).

[53] Moore, A., and Papagiannaki, K. Toward the accurate identification of network applications. In *Proc. of the Passive and Active Measurement Workshop - PAM'2005* (Boston, MA, USA, Apr. 2005).

[54] Ng, T. S. E., and Zhang, H. Predicting Internet network distance with coordinates-based approaches. In *Proc. of the IEEE INFOCOM'2002* (New York, NY, USA, June 2002).

[55] Padmanabhan, V. N., and Subramanian, L. An investigation of geographic mapping techniques for Internet hosts. In *Proc. of the ACM SIGCOMM'2001* (San Diego, CA, USA, Aug. 2001).

[56] Papagiannaki, K., Taft, N., and Diot, C. Impact of flow dynamics on traffic engineering design principles. In *Proc. of the IEEE INFOCOM'2004* (Hong Kong, Mar. 2004).

[57] Papagiannaki, K., Taft, N., and Lakhina, A. A distributed approach to measure IP traffic matrices. In *Proc. of ACM/SIGCOMM Internet Measurement Conference – IMC 2004* (Taormina, Italy, Oct. 2004).

[58] Papagiannaki, K., Taft, N., Zhang, Z., and Diot, C. Long-term forecasting of Internet backbone traffic: Observations and initial models. In *Proc. of the IEEE INFOCOM'2003* (San Francisco, CA, USA, Mar. 2003).

[59] Pásztor, A., and Veitch, D. PC-based precision timing without GPS. In *Proc. of the ACM SIGMETRICS'02* (Los Angeles, CA, USA, June 2002).

[60] Paxon, V., Almes, G., Mahdavi, J., and Mathis, M. Framework for IP performance metrics. *RFC 2330* (May 1998).

[61] Paxson, V. *Measurement and Analysis of End-to-end Internet Dynamics*. PhD thesis, University of California - Berkeley, 1997.

[62] Paxson, V. On calibrating measurements of packet transit times. In *Proc. of the ACM SIGMETRICS'98* (Madison, WI, USA, June 1998).

[63] Paxson, V., and Floyd, S. Wide area traffic: The failure of Poisson modeling. *IEEE/ACM Transactions on Networking 3*, 3 (June 1995), 226–244.

[64] Paxson, V., Mahdavi, J., Adams, A., and Mathis, M. An architecture for large-scale Internet measurement. *IEEE Communications Magazine 36*, 8 (Aug. 1998), 48–54.

[65] Percacci, R., and Vespignani, A. Scale-free behavior of the Internet global performance. *The European Physical Journal B - Condensed Matter 32*, 4 (Apr. 2003), 411–414.

[66] Peterson, L., Anderson, T., Culler, D., and Roscoe, T. A blueprint for introducing disruptive technology into the internet. In *Proc. of the 1st Workshop on Hot Topics in Networks (HotNets-I)* (Princeton, NJ, USA, Oct. 2002). http://www.planet-lab.org.

[67] PIAS, M., CROWCROFT, J., WILBUR, S., HARRIS, T., AND BHATTI, S. Lighthouses for scalable distributed location. In *Proc. of the Second International Workshop on Peer-to-Peer Systems - IPTPS'03* (Berkeley, CA, USA, Feb. 2003).

[68] POSTEL, J. Internet control message protocol. *RFC 792* (Sept. 1981).

[69] PRASAD, R., DOVROLIS, C., MURRAY, M., AND CLAFFY, K. C. Bandwidth estimation: Metrics, measurement tecniques, and tools. *IEEE Network 17*, 6 (Nov. 2003), 27–35.

[70] REKHTER, Y., AND LI, T. A border gateway protocol 4 (bgp-4). *RFC 1771* (Mar. 1995).

[71] SAROIU, S., GUMMADI, P. K., AND GRIBBLE, S. D. Sprobe: A fast technique for measuring bottleneck bandwidth in uncooperative environments, Jan. 2002. `http://sprobe.cs.washington.edu/`.

[72] SHANNON, C., MOORE, D., AND CLAFFY, K. C. Beyond folklore: Observations on fragmented traffic. *IEEE/ACM Transactions on Networking 10*, 6 (Dec. 2002), 709–720.

[73] SHAVITT, Y., AND TANKEL, T. Big-bang simulation for embedding network distances in Euclidean space. In *Proc. of the IEEE INFOCOM'2003* (San Francisco, CA, USA, Mar. 2003).

[74] SHRIRAM, A., MURRAY, M., HYUN, Y., BROWNLEE, N., BROIDO, A., FOMENKOV, M., AND CLAFFY, K. Comparison of public end-to-end bandwidth estimation tools on high-speed links. In *Proc. of the Passive and Active Measurement Workshop - PAM'2005* (Boston, MA, USA, Apr. 2005).

[75] SMITH, P. Cidr report, Mar. 2005. `http://www.cidr-report.org`.

[76] SNOEREN, A. C., PARTRIDGE, C., SANCHEZ, L. A., JONES, C. E., TCHAKOUNTIO, F., SCHWARTZ, B., KENT, S. T., AND STRAYER, W. T. Single-packet IP traceback. *IEEE/ACM Transactions on Networking 10*, 6 (Dec. 2002), 721–734.

[77] SOULE, A., SALAMATIAN, K., EMILION, R., TAFT, N., AND PAPAGIANNAKI, K. Flow classification by histograms or how to go on safari in the Internet. In *Proc. of the ACM SIGMETRICS'04* (New York, NY, USA, June 2004).

[78] SPRING, N., MAHAJAN, R., WETHERALL, D., AND ANDERSON, T. Measuring ISP topologies with Rocketfuel. *IEEE/ACM Transactions on Networking 12*, 1 (Feb. 2004), 2–16.

[79] SPRING, N., WETHERALL, D., AND ANDERSON, T. Reverse-engineering the Internet. In *Proc. of the 2nd Workshop on Hot Topics in Networks (HotNets-II)* (Cambridge, MA, USA, Nov. 2003).

[80] STALANS, L. J. *Reading and Understanding Multivariate Statistics*. APA Press, 1998, ch. Multidimensional Scaling, pp. 137–168.

[81] TANG, L., AND CROVELLA, M. Virtual landmarks for the Internet. In *Proc. of the ACM Internet Measurement Conference 2003* (Miami, FL, USA, Oct. 2003).

[82] TEIXEIRA, R., DUFFIELD, N., REXFORD, J., AND ROUGHAN, M. Traffic matrix reloaded: Impact of routing changes. In *Proc. of the Passive and Active Measurement Workshop - PAM'2005* (Boston, MA, USA, Apr. 2005).

[83] UIJTERWAAL, H., AND ZEKAUSKAS, M. *IP Performance Metrics (IPPM)*, 2003. `http://www.ietf.org/html.charters/ippm-charter.html`.

[84] VARGHESE, G., AND ESTAN, C. The measurement manifesto. *ACM Computer Communication Review 34*, 1 (Jan. 2004), 9–14.

[85] WANG, J., ZHOU, M., AND ZHOU, H. Clock synchronization for Internet measurements: A clustering algorithm. *Computer Networks 45*, 6 (Aug. 2004), 731–741.

[86] ZIVIANI, A., FDIDA, S., DE REZENDE, J. F., AND DUARTE, O. C. M. B. Toward a measurement-based geographic location service. In *Proc. of the Passive and Active Measurement Workshop - PAM'2004* (Antibes Juan-les-Pins, France, Apr. 2004), Lecture Notes in Computer Science (LNCS) 3015, pp. 43–52.

[87] ZIVIANI, A., FDIDA, S., DE REZENDE, J. F., AND DUARTE, O. C. M. B. Improving the accuracy of measurement-based geographic location of Internet hosts. *Computer Networks 47*, 4 (Mar. 2005), 503–523.

[88] ZIVIANI, A., AND SCHULZE, B. Measurement middleware service for grid computing. In *Poster in the 2nd International Workshop on Middleware for Grid Computing - MGC 2004* (Toronto, Canada, Oct. 2004).

Artur Ziviani received the B.Sc. degree in Electronics Engineering in 1998 and the M.Sc. degree in Electrical Engineering (with emphasis in Computer Networking) in 1999, both from the Federal University of Rio de Janeiro (UFRJ), Brazil. In December 2003, he received the Ph.D. degree in Computer Science from the University Pierre et Marie Curie (Paris 6), Paris, France, where he has also been a lecturer during the 2003-2004 academic year. Since September 2004, he is with the National Laboratory for Scientific Computing (LNCC), a research unit of the Brazilian Ministry of Science and Technology, located in Petrópolis, Brazil. His major research interests are quality of service (QoS) provisioning, mobile and wireless computing, and Internet measurements.