GlobeTraff: a traffic workload generator for the performance evaluation of future Internet architectures

Konstantinos V. Katsaros, George Xylomenos and George C. Polyzos
Mobile Multimedia Laboratory
Department of Informatics
Athens University of Economics and Business
Patision 76, Athens 104 34, Greece
E-mail: ntinos@aueb.gr, xgeorge@aueb.gr and polyzos@aueb.gr

***Abstract*—Having identified important limitations of the current Internet architecture, several research initiatives have engaged in the design of new architectures for the Internet of the Future. New features and protocols are designed, and in many cases, a *clean slate* approach is followed, advocating the replacement of almost the entire current protocol stack. In order to reliably evaluate the performance of such emerging protocols and architectures, we need a realistic evaluation framework reflecting current and forecasted traffic patterns. This is especially important for the emerging information-centric paradigm, where in-network caching plays an important role and network performance is heavily dependent on traffic characteristics. To this end, we have designed and implemented *GlobeTraff*, a traffic workload generator for the creation of synthetic traffic mixes. GlobeTraff supports several application traffic types, generated via models in the recent research literature, allowing the detailed parameterization of the respective models and the composition of the resulting traffic mix.**

***Index Terms*—workload; simulation; emulation; caching; performance evaluation**

I. Introduction

The tremendous success of the Internet is amply demonstrated by its exponential growth and its deep immersion into almost every aspect of social and economic activity. However, the networking research community has realized that the current Internet architecture exhibits significant limitations (e.g., mobility support, scalable content distribution) which hinder its evolution and the efficient accommodation of evolving user needs. Taking a step away from addressing these issues with patches or application specific solutions, the networking research community has engaged in an effort to define alternative networking paradigms [1]. In many cases these efforts even consider the replacement of (almost) the entire TCP/IP protocol stack, yielding *clean slate* designs (e.g. [2]).

Research on new networking architectures calls for the careful design of proper performance evaluation mechanisms, able to support a better understanding of the new functionality, assess the benefits stemming from the new approaches and enable the fine tuning of the resulting protocols. In this context, the traffic patterns these new architectures are called to serve play a vital role. To this end, past research on traffic modeling has yielded significant results for various types of traffic e.g., Web, P2P, Video. In all these cases, the derived models have served as a valuable input in assessing the performance characteristics of application-specific traffic types. However, the advent of clean slate approaches for the future Internet changes this application-specific point of view. The emerging protocols refer to the Internet as a whole, therefore their design must adhere to the needs of *all* anticipated application types and the corresponding traffic patterns.

So far, ICN-oriented performance evaluation efforts either focus on application specific traffic models (e.g., [3]) or use simplified traffic mixes (e.g., [4]). In the former case, the resulting traffic workloads provide a partial view of network traffic, neglecting the impact of traffic diversity in the Internet e.g., Web traffic and/or P2P traffic. This approach is sufficient to assess specific performance issues, but is unable to characterize the overall performance of an architecture. In the latter case, the traffic mixes employed do not reflect the detailed characteristics of each participating traffic type, often overlooking certain types of traffic or oversimplifying important aspects, such as the exact shape of the popularity distribution and the temporal evolution of popularity.

The critical role of realistic workload models in the performance evaluation of new networking architectures becomes more evident in the case of *Information-Centric Networking* (ICN). The ICN paradigm reflects the fact that current end users primarily focus on the information itself, rather than on the end hosts providing it. In ICN architectures information is named and the overall operation of the networks follows the publish/subscribe paradigm i.e., information providers (publishers) denote the availability of information to the network and content consumers (subscriber) issue requests to the network in order to indicate the desired piece of information. It is then the responsibility of the network to act as a mediator, by locating and delivering the content. A significant characteristic of the ICN paradigm is that, by naming information, it allows the deployment of in-network caching, which, contrary to earlier approaches, operates at the network level and is oblivious to the applications generating the data. In most cases, these caches operate at a packet level (e.g., [4]), but other granularities have also been investigated (e.g.,[3]).

In this paper we present GlobeTraff, a synthetic traffic workload generator that enables the creation of traffic mixes reflecting the current characteristics of Internet traffic. GlobeTraff extends the functionality of an earlier tool focused on web traffic [5], [6], by further providing support for P2P and video traffic. GlobeTraff also allows the creation of various traffic mixes, thus enabling the investigation of several scenarios/trends in the composition of Internet traffic. In the remainder of this paper, we first outline the traffic model aspects that affect network performance evaluation (Section II), paying particular attention to the characteristic case of ICN and discussing how the presented traffic model aspects may affect the perceived performance of in-network caching schemes. We then present in detail the set of traffic models used in GlobeTraff and their parameters (Section III) and finally provide an overview of our software tool (Section IV).

II. Traffic characteristics

A. Traffic mixture

A significant challenge in characterizing/modeling traffic on the Internet relates to the identification of the applications contributing to the observed traffic (e.g., Web, P2P), as well as to the exact amount of traffic they contribute to the overall traffic mixture. This is especially important in the case of in-network caching where different types of content, with different popularity characteristics and object sizes (see next subsections) compete for the same caching space. The composition of the overall traffic mixture evolves as new applications, protocols and services emerge. Measurement based studies provide only snapshots of this constant evolutionary process, constituting a first step in the effort of generating synthetic Internet traffic workloads. This makes important the flexibility to express various traffic mixtures in a synthetic workload.

B. Popularity characteristics

One of the most important aspects of Internet traffic relates to the popularity of the delivered content. These characteristics strongly relate to the ability to reduce redundant transmissions of the same content, thus reducing traffic, content provider load and delivery times. Popular content favors the deployment of caching mechanisms, since data is transferred once from its source and subsequent requests are served from the cache. Similarly, multiple requests for the same content can possibly be served via (application level) multicast. However, in both cases, the efficiency of the deployed mechanisms heavily depends on the spatio-temporal characteristics of content popularity, as discussed in the following.

1) Popularity distribution: The popularity of an information item is usually expressed by the fraction of requests it receives over all requests submitted for the entire set of considered information items, or simply by the number of requests for it. In most application types considered in this work, the popularity distribution follows some power law, indicating that a small number of information items triggers a major portion of the entire set of requests. The exact shape of the power law popularity distribution directly impacts the performance of the underlying protocols. For instance, highly skewed popularity distributions (e.g., a Zipf-like distribution with a high slope value) favor the deployment of caching schemes, since caching a very small set of information items can dramatically increase the cache hit ratio.

2) Temporal locality: The popularity distribution expresses the number of requests submitted for each information item participating into a certain workload. However, they do not describe how these requests are distributed in time. This aspect is of primary importance when considering the performance of caching schemes since the ordering of the requests obviously affects the contents of a cache. For example, with a *Least Frequently Used* (LFU) cache replacement policy, if all requests for a certain item are submitted close in time, the item is unlikely to be evicted from the cache, even by a (globally) more popular item whose requests are more evenly distributed in time. The temporal ordering of requests gains even more importance when considering workloads consisting of various applications, all competing for the same cache space.

3) Spatial locality: Spatial locality refers to the distribution of the requests across the network. The importance of spatial locality relates to the ability to avoid redundant traffic in the network. If requests are highly localized in some area of the entire network, then similar requests can be more efficiently served with mechanisms such as caching and/or multicast i.e., the concentration of similar requests in a limited area of the network allows increasing the perceived cache hit ratios at caches in the area and/or the traffic savings from the use of multicast. This aspect of popularity is not reflected by GlobeTraff, as it strongly relates to the network topology used in the performance evaluation. However, we consider it of primary importance for the proper evaluation of network protocols and/or distributed caching schemes.

C. Object sizes

Another important aspect of the generated workload is that of the distribution of the size of the information items described. Obviously, the size of each item affects the overall traffic load in the network and thus has a direct impact on the performance of the evaluated protocols/architectures. In addition, if we consider an application-oblivious caching scheme, the different information item sizes of each application may result in a highly inefficient cache space utilization.

III. Models Implemented

The GlobeTraff tool takes into account the traffic characteristics discussed above in order to produce synthetic workloads that can be used in the evaluation of new network protocols, architectures and/or caching schemes. Currently GlobeTraff supports models found in the literature for the cases of Web, P2P and Video workloads. This section describes the specific models adopted in GlobeTraff and their parameters, while the next section provides an overview of the software itself. It should be noted that GlobeTraff allows the user to select the desired values for all considered model parameters, as well as for the proportions of each traffic in the overall traffic mix.

A. Traffic mixture

In our work we have resorted to two recent measurement studies that investigate the composition of Internet traffic [7], [8]. In [7] a large scale measurement study was performed, with the purpose of studying the traffic crossing inter-domain links. The results indicate the dominance of Web traffic, amounting to 52% over all measured traffic. However, *Deep Packet Inspection* (DPI) techniques reveal that 25-40% of all HTTP traffic actually carries video traffic. Results from DPI techniques also reveal the difficulty in correctly identifying the application type in the case of P2P traffic: mapping observed port numbers to well known applications shows P2P traffic constituting only 0.85% of overall traffic, while DPI raises this percentage to 18.32% [7]. Relevant studies on a large ISP show the percentage of P2P traffic ranging from 17 to 19% of overall traffic [8]. Table I provides an overview of these figures. The *Other* traffic type denotes traffic that cannot be classified in any of the first three application categories, and consists of unclassified traffic and traffic heavily fragmented into several applications (e.g., 0.17% DNS traffic).

Traffic Type	Percent of Total Traffic
Web	31.2-39%
P2P	17-19%
Video	13- 20.8%
Other	29-31%

TABLE I
TRAFFIC MIX STATISTICS.

B. Web traffic

1) Popularity distribution: It has been shown that the popularity of web information items is well modeled by a Zipf-like distribution of the following form:

$$p(i) = \frac{K}{i^a} \qquad (1)$$

where i is the popularity rank, N is the total number of considered information items, a is the slope of the distribution and $K = 1/\sum_{i=1}^{N} \frac{1}{i^a}$. Of particular importance is the value of a which determines the skewness of the distribution and, in effect, the number of the items with higher popularity values. Several measurement based studies have yielded a values in the range of 0.64-0.84 [9], [10].

2) Temporal Locality: The ordering of the requests has been modeled with an LRU stack model [5]. Based on this ordering, the exact timing of the requests is then described by an exponential distribution [11].

3) Object Sizes: It has been shown that the distribution of web object sizes is well modeled by the concatenation of the Lognormal (body) and Pareto (tail) distributions [12], [13].

C. P2P traffic

1) Popularity distribution: It has been shown that in typical P2P applications, information item popularity is expressed by a Mandelbrot-Zipf distribution of the following form [14]:

$$p(i) = \frac{K}{(i+q)^a} \qquad (2)$$

where q is a *plateau* factor, with typical values ranging from 5 to 100, yielding a flatter distribution head i.e., more items exhibit high popularity values than in the corresponding Zipf-like distribution (where $q = 0$). The other parameters (a, i and K) are defined as in the Zipf distribution (see above).

2) Temporal Locality: Here we adopt the models derived in [15] for the BitTorrent application, as BitTorrent has been shown to be the most prevalent P2P application [8]. Based on application traces, it was shown that the average arrival rate of new torrents (and hence new items in the workload) is equal to 0.9454 torrents per hour. In the same study it was also observed that the arrival rate of peers in a swarm follows an exponential decreasing rule with time t:

$$\lambda(t) = \lambda_0 e^{-\frac{t}{\tau}},$$

where λ_0 is the initial arrival rate when the torrent starts and τ expresses the torrent's popularity. Based on this distribution, it can be shown that $N_{all} = \lambda_0 \tau$, where N_{all} is the total torrent population size. We derive the average values for λ_0 and τ from [15] (87.74 and 1.16 respectively). Note that the derived τ value expresses the popularity of the average P2P object. We therefore adjust its value for each specific object, according to its popularity. Based on these models, GlobeTraff initially generates a random ordering of swarm birth events i.e., the timing of the first request for each P2P object is determined. Then, for each swarm the corresponding τ value is calculated, and the above exponential decay rule is applied to determine the exact timing of each subsequent request in the swarm.

3) Object Sizes: It has been shown that file sizes vary widely in the BitTorrent application[14], [16]. To preserve realism, we have resorted to sampling of BitTorret traces [16] i.e., each P2P information object size is randomly selected from a trace of real BitTorrent files. The option to use a fixed value is also provided by GlobeTraff.

D. Video traffic

Regarding, video traffic, we consider the most popular service in the Internet i.e., YouTube.

1) Popularity distribution: Studies have shown that a Zipf-like distribution does not adequately model the popularity of YouTube videos [17]. Instead, it was shown that the Weibull distribution (k=0.513, λ=6010) or the Gamma distribution (k=0.372, θ=23910) can better model the popularity of YouTube videos.

2) Temporal Locality: To the best of our knowledge there are no available studies modeling the temporal evolution of YouTube video popularity. GlobeTraff randomly distributes all video requests throughout the total duration of the generated synthetic workload; this duration is determined by the P2P traffic requests.

3) Object Sizes: Similarly to P2P objects, YouTube videos vary widely in length. In [18] it was shown that YouTube video length can be well modeled by a concatenated normal distribution. Since the vast majority of YouTube videos have a bit-rate close to 330 Kbps [18], there is a direct mapping between video length and file size, hence we use the concatenated normal distribution to model video file sizes.

E. Other traffic

1) Popularity distribution: Based on the fact that a Zipf-like distribution can well describe object popularity for most application types, we also use it to model the popularity of other traffic objects.

2) Temporal Locality: Since this is other traffic, we follow the same approach as with the Video traffic type, with the additional option of adopting the Web based model of exponential request inter-arrival times. In the latter case, the entire set of requests (to all objects in this category) is shuffled and follows an exponential distribution of inter-arrival times, adjusted to fit the entire duration of the synthetic workload.

3) Object Sizes: The file sizes for the other traffic can also vary widely. GlobeTraff allows the user to set the size of the objects in this category. In order to determine a default value and further provide an indication to the GlobeTraff user, we first note that besides serving human communication, the future Internet is also expected to be an "Internet of Things" enabling machine to machine communications. Thus, we expect huge amounts of small pieces of information to be produced and consumed over the infrastructure. Therefore, compared to the other traffic types, we select a relatively small object size for this type of traffic i.e., 5KB.

IV. The GlobeTraff tool

A. Implementation

The implementation of GlobeTraff is based on the ProWGen trace generation tool [5], written in C++. We used ProWGen for the generation of the Web traffic workload and then heavily extended it by further implementing the models described above for the P2P, Video and Other workload types. Moreover, we implemented a Java-based *graphical user interface* (GUI) that facilitates the parameterization of all considered models, so as to ease the generation of synthetic workloads.

B. Usage

GlobeTraff allows the user to easily parameterize all considered models. As shown in the sample screenshot of Figure 1, the GlobeTraff front-end allows the user to define the exact composition of the generated traffic mix, as well as its total size. Then, for each traffic type the respective models (see Section II) can be parameterized. For instance, Figure 1 shows the available options for the generation of the P2P synthetic workload. The user can change the default parameter values for the Mandelbrot-Zipf popularity distribution, as well as the exponential decay law describing peer arrivals. The user may also select whether to sample P2P file sizes from real traces or use a fixed value for them. The remainder of GlobeTraff's GUI options is not presented due to length limitations. The source code of the tool is available for download [19].

1) Output: Following the ProWGen convention, the generated workload is described in two files, so as to allow tools compatible with ProWGen to directly take advtantage of GlobeTraff. The first one describes the objects appearing in the workload, as shown in Table II. A distinct line is created for each object, containing a unique identifier (ID) for the object, its popularity expressed as a total number of requests, its size in bytes and the application type it belongs to (1: Web, 2: P2P, 3: Video, 4: Other).

Item ID	Popularity	Size (Bytes)	Application Type
0	34	8145	1
1	18	9033	1
2	12	8751	1
3	12	8145	1

TABLE II
OBJECTS DESCRIBED IN THE WORKLOAD.

The second file describes the workload with respect to time i.e., each line refers to a single request and denotes the time the request is submitted, the ID of the Object that the request refers to, and the size of the object as shown in Table III. GlobeTraff produces a separate pair of files for each traffic type, as well as a concatenated set describing the entire workload mixture.

Time	Item ID	Size (Bytes)
0.018314	0	8145
0.171501	2	8751
1.374289	1	9033
4.240712	1	9033

TABLE III
TIMING OF REQUESTS FOR OBJECTS.

V. Conclusions and future work

In this paper we argued about the importance of employing realistic traffic models for the proper performance evaluation of new network protocols and architectures and further pointed out the importance of using non application-specific traffic mixes for ICN architectures. To this end, we outlined several key aspects of traffic models and discussed their importance and impact on the performance evaluation process. Based on these observations, we presented GlobeTraff, a highly parameterizable tool for the generation of realistic, synthetic workload mixes. GlobeTraff provides support for numerous traffic models described in the literature.

While the traces produced by GlobeTraff are useful for all kinds of Internet performance evaluation studies, we consider the tool to be especially useful for studies of future Internet architectures and protocols, as the ability to parameterize both the traffic mix and the individual traffic type characteristics, allows researchers to make their own projections about future Internet traffic patterns and generate the appropriate traces. Moreover, GlobeTraff has been implemented in a modular way facilitating the implementation of new workload models.

Regarding future work, we have identified several key limitations of currently available models. Most importantly, the temporal evolution of the popularity characteristics of video items has not been thoroughly investigated. Considering the increasing participation of video traffic in the Internet traffic mix, this issue is considered of particular importance and is a prime area for further study.

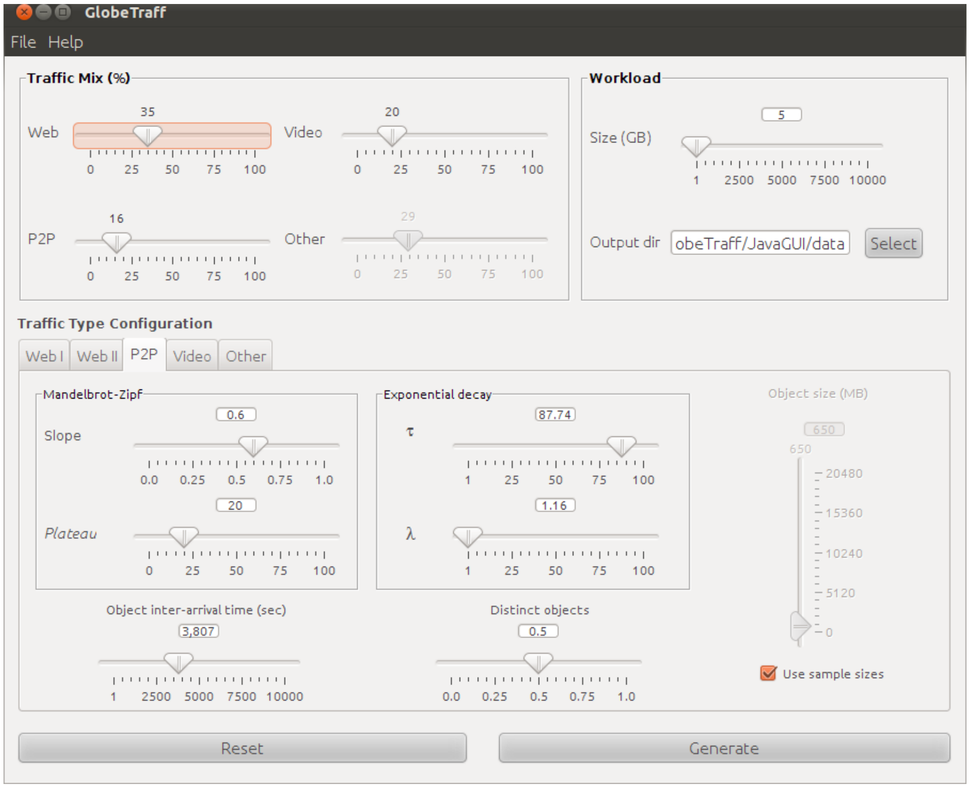

Fig. 1. GlobeTraff GUI.

Acknowledgment

The work reported in this paper was supported by the FP7 ICT project PURSUIT, under contract ICT-2010-257217.

References

[1] J. Pan, S. Paul, and R. Jain, "A survey of the research on future Internet architectures," *Communications Magazine, IEEE*, vol. 49, no. 7, pp. 26–36, 2011.

[2] PURSUIT Project, *PURSUIT Project Home Page*, http://www.fp7-pursuit.eu, 2011.

[3] K. Katsaros, G. Xylomenos, and G. C. Polyzos, "MultiCache: An overlay architecture for information-centric networking," *Computer Networks*, vol. 55, pp. 936–947, 2011.

[4] G. Carofiglio, V. Gehlen, and D. Perino, "Experimental evaluation of storage management in content-centric networking," in *ICC*, 2011.

[5] M. Busari and C. Williamson, "ProWGen: a synthetic workload generation tool for simulation evaluation of web proxy caches," *Computer Networks*, vol. 38, no. 6, pp. 779–794, 2002.

[6] N. Markatchev and C. Williamson, "WebTraff: A GUI for Web Proxy Cache Workload Modeling and Analysis," in *IEEE MASCOTS*, 2002, pp. 356–363.

[7] C. Labovitz, S. Iekel-Johnson, D. McPherson, J. Oberheide, and F. Jahanian, "Internet inter-domain traffic," in *ACM SIGCOMM*. ACM, 2010, pp. 75–86.

[8] G. Maier, A. Feldmann, V. Paxson, and M. Allman, "On dominant characteristics of residential broadband internet traffic," in *ACM IMC*, 2009, pp. 90–102.

[9] L. Breslau, P. Cao, L. Fan, G. Phillips, and S. Shenker, "Web Caching and Zipf-like Distributions: Evidence and Implications," in *IEEE INFOCOM*, 1999, pp. 126–134.

[10] A. Mahanti, C. Williamson, and D. Eager, "Traffic analysis of a web proxy caching hierarchy," *IEEE Network*, vol. 14, no. 3, pp. 16 –23, 2000.

[11] M. F. Arlitt and C. L. Williamson, "Internet web servers: workload characterization and performance implications," *IEEE/ACM Transactions on Networking*, vol. 5, pp. 631–645, 1997.

[12] P. Barford and M. Crovella, "Generating representative web workloads for network and server performance evaluation," in *ACM SIGMETRICS/PERFORMANCE*, 1998, pp. 151–160.

[13] P. Barford, A. Bestavros, A. Bradley, and M. Crovella, "Changes in web client access patterns: Characteristics and caching implications," *World Wide Web*, vol. 2, pp. 15–28, 1999.

[14] M. Hefeeda and O. Saleh, "Traffic Modeling and Proportional Partial Caching for Peer-to-Peer Systems," *IEEE/ACM Transactions on Networking*, vol. 16, no. 6, pp. 1447–1460, 2008.

[15] L. Guo, S. Chen, Z. Xiao, E. Tan, X. Ding, and X. Zhang, "A performance study of BitTorrent-like peer-to-peer systems," *IEEE Journal on Selected Areas in Communication*, vol. 25, no. 1, pp. 155–169, 2007.

[16] A. Bellissimo, B. N. Levine, and P. Shenoy, "Exploring the use of BitTorrent as the basis for a large trace repository," University of Massachusetts Amherst, Tech. Rep., 2004.

[17] X. Cheng, C. Dale, and J. Liu, "Statistics and social network of youtube videos," in *IEEE IWQoS*. IEEE, 2008, pp. 229–238.

[18] ——, "Understanding the Characteristics of Internet Short Video Sharing: YouTube as a Case Study," *CoRR*, vol. abs/0707.3670, 2007.

[19] K. Katsaros. (2012, Jan) Globetraff. [Online]. Available: http://www.mm.aueb.gr/~katsaros/GlobeTraff.tar.gz