A Prediction Criterion for Selecting Popular Destinations

A. Fonte*†, M. Curado*, E. Monteiro*
* University of Coimbra, Polo II, Pinhal de Marrocos, 3030-290, Coimbra, Portugal
Email: {afonte,edmundo,marilia}@dei.uc.pt
†Polytechnic Institute of Castelo Branco, ESTCB, Av. do Empresario, 6000-767, Castelo Branco, Portugal

Abstract—**Inter-domain traffic engineering is being increasingly integrated in network management systems. Yet, scalability issues of traffic engineering may limit its applicability to large networks due to excessive computational costs and overhead of routing changes. One documented means of reducing these overheads is to focus the routing optimizations on the paths to popular destinations, and thus be able to shift a large volume of traffic with a small number of path switches, rather than shifting the traffic for all the prefixes. However, there is a lack of simple and pragmatic methods for selecting popular destinations. This paper seeks to address this problem by introducing a practical criterion to define a threshold to categorize the popularity of the traffic. This is based on the analysis of the errors of commonly used predictors for traffic tracking. Our results show that by applying this criterion, we can reduce the number of target prefixes to a small fraction of the total number, while ensuring stability in the traffic engineering, which results from an effective predictability of the traffic headed to these prefixes.**

I. Introduction

Inter-domain Traffic Engineering (TE) has become an important part of today's network management systems to provide best inter-domain routing, while ensuring a high level of global network performance [1], [2]. This is particularly important to generate cost savings and to face several challenges posed by the competitive Internet market, such as the provision of value-added services.

The inputs to the inter-domain traffic engineering process are the traffic demands over the network, the available egress point choices for the destinations (provided by BGP (Border Gateway Protocol) [3]) and the egress point capacities. The output is the inter-domain routing so as the traffic objectives can be satisfied. Lastly, the optimal routing is translated to a careful tuning of the BGP routes attributes [4].

One of the common optimization problems that has to be dealt with by the network managers is, thus, the Egress Router Selection (ERS) [5]: *how should the traffic demands be assigned to multiple egress points, to ensure that the transit network's traffic objectives are fulfilled (e.g., minimizing the maximum link utilization (min-MLU) or Load-balancing (LB))*? Studies, such as [6], [7], provide some instances of algorithms which can solve the problem of ERS efficiently.

The biggest concern with ERS is that with regard to the number of traffic flows, objectives and egress point choices, the task of optimizing the routing can become computationally hard and entails a large number of route changes. One effective means of reducing these overheads is to focus the traffic engineering optimization on popular destinations (a.k.a. top receivers), and thus being able to shift a large volume of traffic with a small number of path switches, rather than shifting the traffic for all the prefixes [1].

However, there are two major issues that have to be tackled; these are the selection of popular destinations, and the prediction of the amounts of traffic that will be sent to each destination during the next time interval. The selection of popular destinations is of critical importance to alleviate the complexity of the traffic engineering process, which has an inherent problem of lacking scalability. Tracking the traffic accurately is a significant issue because the network traffic is bursty [8], which means that weak predictions may lead to spurious traffic changes or congestion over egress links or downstream Internet Service Providers (ISP).

This paper correlates both issues and sets out a practical criterion to define a threshold for traffic volumes, the value of which is used to categorize the popularity of the traffic destinations. This is based on the analysis of errors of common predictors for traffic tracking, namely (the simple) Last-Value (LV), the Moving Average (MA), and the Low-pass Exponential Moving Average (LpEMA). Our results show that by applying this criterion, we can reduce the number of target prefixes to a small fraction of the total number, while ensuring stability in the traffic engineering, which results from an effective predictability of the traffic headed toward these prefixes.

The rest of this paper is organized as follows. Section II outlines the whole inter-domain traffic engineering process, describes the concept of Zipf's Law and its significance for traffic engineering, and states our problem. Sections III describes the data trace and analyzes the performance of different traffic predictors to postulate a criterion for selecting popular destinations. Finally, Section IV concludes this paper.

II. Background and Statement of the Problem

This section sets out by describing the traffic engineering process. We then underline the importance of the consistency of traffic demands with the Zipf's law for the traffic engineering. Lastly, we describe the problem of selecting the popular destinations, and present our proposal.

This work is partly supported by the FCT User Centric Routing (UCR) Project (PTDC/EEA-TEL/103637/2008).

A. Basics of Inter-domain Traffic Engineering

The goal of inter-domain traffic engineering is to optimize inter-domain routing, subject to a given traffic objective and the constraints of the network. For the sake of illustration, we consider as traffic objective, the minimization of the maximum link utilization (MLU) of the egress links. A transit network to be optimized is represented by a set of ingress points I and a set of egress points E.

The inputs to the TE algorithm are the incoming traffic demands (TD) over the transit network, the available egress point choices for each reachable prefix destination $p \in P$ (provided by BGP) and the egress point capacities C, where each entry $c(e)$ is the capacity at egress point $e \in E$. The predicted TD are represented by the matrix $D = \{d_{ip} \mid i \in I, p \in P\}$, where each entry d_{ip} is the demand for the ingress point i - destination p pair. The set of egress point capacities is represented as C, where each entry $c(e)$ is the capacity at the egress point e. An inter-domain routing is represented by ϵ, where each entry $\epsilon_{iep} \in \{0, 1\}$ is an indicator function that tells whether the d_{ip} is assigned to the egress point e. The output of the inter-domain traffic engineering process is the set of optimal routes to achieve the optimal traffic mixture. In turn, these results are translated to a careful tuning of BGP routes attributes. For the interested reader, the set of techniques (e.g., `LOCAL-PREFs` tuning) that can be used for egress traffic control are described in [4].

The TE problem is known as the egress router selection (ERS): *how to assign each entry of traffic demands d_{ip} to an egress point e, so as to optimize a certain traffic objective*. The traffic objective that is encoded in the ERS problem is to minimize the whole MLU (min-MLU) at egress links. To be more specific, we introduce the definitions 2.1 and 2.2.

Definition 2.1: *The link utilization of e for a routing ϵ is defined as the traffic to capacity ratio as shown in* (1).

$$U_e = \sum_i \sum_p \frac{\epsilon_{iep} d_{ip}}{c(e)}, \epsilon\ is\ a\ routing. \tag{1}$$

Definition 2.2: *An optimal inter-domain routing for a given D, is the routing that minimizes the maximum link utilization (min-MLU), as shown in (2), where OU is the optimal utilization.*

$$OU = min\ max\ U_e, \forall e \in E. \tag{2}$$

The traffic objective min-MLU is subject to constraints (3) and (4). The capacity constraint (3) ensures that the total resource requirements of the traffic flows assigned to each egress point do not exceed the available/contracted capacity. The assignment constraint (4) guarantees that each traffic flow is assigned to exactly one egress point e.

$$\sum_i \sum_p \epsilon_{iep} d_{ip} \leq c(e), \forall e \in E \tag{3}$$

$$with,\ \sum_e \epsilon_{iep} = 1, \forall i \in I \tag{4}$$

B. Zipf's Law and its Significance for Traffic Engineering

Zipf's law is an empirical law introduced to describe the popularity of words in terms of rank and their frequency in use [9]. It states that if f_i is the frequency of the word i and r_i its rank order, then $f_i \simeq \frac{1}{r_i}$. This implies that the n-th word is used twice as often as the $2n$-th word. Visually Zipf's law can be easily observed by plotting the data set on a log-log scale, with the axes being $x = log(rank\, order)$ and $y = log(frequency)$. If the plot is (almost) linear, we say that the data set is consistent with Zipf's law. When using Zipf's law in general contexts, such as the popularity of web pages or traffic destinations as in this paper, it can be re-formulated as: if f_i as a function of r_i is consistent with a power-law distribution it is referred to as Zipf's-like.

A typical backbone network has routes for more than 150000 prefixes [1]; as a result optimizing the routing to accommodate all the traffic demands that were detected by network monitors, can become computationally hard and entail a large number of route changes. The extrapolation of this concept to the field of traffic engineering is therefore significant.

When traffic demands of Internet traces (in terms of traffic volumes) are consistent with the Zipf's-like distribution, it implies that a small fraction of prefixes (i.e., about 5 - 10%, though in our trace about 2% of the total number of prefixes (see Section III)) of the total number of receivers contribute to most of the total volume of traffic. In practice, several studies have identified this phenomenon in Internet traffic [10], [11] which means that it is not necessary to consider all the destinations in the routing optimization process. It will be enough to change `LOCAL-PREFs` for a small handful of popular destinations, that would move a large fraction of traffic from one egress point to another, while minimizing the number of path changes. Moreover, these destinations tend to have more stable traffic volumes [11].

The process of verifying the consistency of traffic demands with Zipf's law proceeds as follows. First, the traffic demands are acquired from measurements taken from a network, i.e., the total amount of traffic sent to each destination is recorded during an interval of time [10], [12]. Subsequently, the flow rankings are computed on the basis of the contribution that each flow makes to the total traffic volume.

A flow ranking R is an ordered set of flows as shown in definition 2.3. The consistency of traffic demands with a Zipf's-like distribution is evaluated according to definition 2.4 or visually by plotting the data on a log-log scale.

Definition 2.3: *Given a set of k flows represented by a vector $F = (f_1, ..., f_k)$, R is a ranking of flows iff $\forall f_i, f_j \in F$ $r_i \succeq r_j \Leftrightarrow V_{norm}(f_j) \leq V_{norm}(f_i)$, where r_k is the rank of flow f_k, $V_{norm}(f_k)$ represents the traffic in terms of bytes for the flow f_k normalized by total volume of traffic.*

Definition 2.4: *Given a set of k flows $F = (f_1, ..., f_k)$, and a ranking R, these are consistent with a Zipf's-like distribution, iff $\forall f_i$ the corresponding traffic volume $V(i)$ satisfies the relation $V(i) = c.r_i^{-\alpha}, \alpha > 0, c = constant$.*

C. Problem Statement, and Challenge

As a network/ISP grows, a commonly-used traffic engineering practice is to select the popular destinations of the traffic and optimize the traffic performance mainly for them, as previously described. However, there remains an important issue to address which is, *how to define a proper value for a threshold (said T) that splits the whole set of destinations into popular and non-popular destinations*. This means that after a ranking R has been computed, both the flows and destinations, whose contribution to the total traffic volume is below some specified threshold T, i.e., $V_{norm}(.) < T$, should be pruned from the optimization process of the network. The remaining flows are considered as top receivers.

Unfortunately, the optimal value for T is hard to find as it depends on several factors, such as the trade-off between the overhead on the traffic engineering algorithm, the overhead of routing changes and the degree of routing control. The purpose of this paper is not to deal with the issue of showing formally what is the best choice for threshold T, but rather to provide a practical method for setting the threshold T.

In this paper, we postulate that the threshold T can be empirically found through the analysis of the errors of the predictors that are used for traffic tracking. A prediction error is defined as in definition 2.5.

Definition 2.5: *Given a flow a prediction error is the difference between the value of the estimated traffic volume e_i for the next time slot i, and the real value V_i, i.e., $error_i$=$|e_i - V_i|$.*

III. Prediction Criterion for Selecting Popular Destinations

Our thesis is that the analysis of the whole behavior of the prediction error in tracking traffic leads to a practical criterion for selecting popular destinations. In this section we, thus, study the performance of different predictors (LV, MA and LpEMA). In evaluating the performance of the predictors, their mean error was computed and analyzed.

A. Data Trace

We use one set of data traces that was collected at the GÉANT pan-European academic network in 2005 [13]. This was conducted by means of Cisco's Netflow [14]. The Netflow measurements were carried out over a period of two weeks, as shown in Fig. 1. Each sample in the figure represents the amount of bytes seen during each interval of 15 minutes, multiplied by 1000, since the Netflow sampling was performed at a rate of 1/1000. After this, we divided the traffic into demands. A demand represents a given amount of traffic (in bytes) from GÉANT's users to a destination.

B. Tracking Traffic

In this study, we track traffic with three distinct predictors which vary in degrees of complexity: a very simple predictor, the Last Value (LV), the classical Moving Average (MA) , and an adaptive but more complex predictor, the LpEMA (Low pass Exponential Moving Average), defined below.

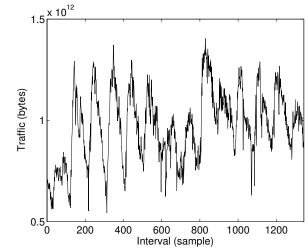

Figure 1. GÉANT Traffic: (left) Traffic evolution from Sunday (05-07-10) to Saturday (05-07-16).

LV is the basic predictor. The actual traffic estimate e_i is equal to the traffic volume V_i measured in the last time slot i, i.e., $e_i = V_i$.

MA is another very simple predictor. The actual traffic estimate e_i is equal to the arithmetic mean of the traffic volumes measured in the last n time slots, i.e., $e_i = \frac{\sum_{k=0}^{n-1} V_{i-k}}{n}$. This predictor requires a set of n traffic volume measures, and thus a window of size n. The biggest problem of this predictor is how to find the right size for the sliding window. Very large windows result in smoothing too much the real traffic changes. In contrast, small windows result in that fast traffic changes are not suppressed.

LpEMA is an extension of the classic EMA predictor, and is a more complex predictor than LV and MA. To compute the actual metric estimate e_i, the LpEMA combines the previous estimate e_{i-1} with the actual traffic volume measure V_i using an adaptive Exponential Moving Average, as shown in (5), where α_i is an adaptive exponential weight, which is calculated by using the classical formula for low pass filter, m_i is the gradient between two metric samples (i.e., $\frac{V_i - V_{i-1}}{t_i - t_{i-1}}$), and m_{norm} is the normative gradient calculated over a given time window (e.g., 10 times the interval $t_i - t_{i-1}$). In contrast to the original EMA, it makes use of an adaptive exponential weight α, since with large weights the estimation follows the measurement exactly, but does not suppress fast traffic changes, whereas with small weights, the traffic changes are suppressed but the estimation follows the real changes too slowly [15].

$$\begin{cases} e_i = (1 - \alpha_i) e_{i-1} + \alpha_i V_i \\ \alpha_i = \alpha_{max} \frac{1}{1 + \frac{|m_i|}{m_{norm}}} \end{cases} \tag{5}$$

More complex predictors could also be employed in our study, such as Auto-Regressive Integrated Moving-Average (ARIMA) (that combines linearly past traffic volumes and/or errors) [16] and Neuronal Networks (NN) (here the basic idea is to train a NN with past traffic volumes to predict future values) [17]. We only employed LV, MA and LpEMA in the belief that there is no advantage in using complex predictors given the fact that the performance achieved is almost the same as with the simpler predictors [18], [19].

C. Analysis of Prediction Errors

We start by verifying the consistency of the GÉANT's data trace with Zipf's law using the procedure described in Section

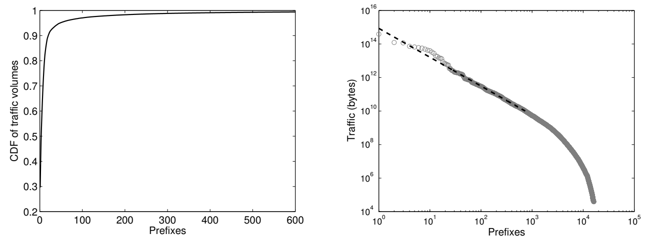

Figure 2. GÉANT Traffic: (left) Cumulative Distribution Function of traffic volumes; (right) Zipf's law fitting giving $\alpha = 1.5818$ and $c = 9.6634E^{10}$.

II-B. After the completion of the process, the right-hand graph of Fig. 2 shows that the trace is roughly consistent with the Zipf distribution with $\alpha = 1.5818$ and $c = 9.6634E^{10}$. This observation implies that a small fraction of the prefixes of the GÉANT data trace contribute with most of the total volume of traffic; this allows us to proceed with the analysis and select GÉANT's popular destinations. It should be noted that to fit the Zipf distribution, we first fit the traffic volumes with a Power-law distribution by using the method described in [20]. Then, we map the Power-law distribution that has been found in a Zipf distribution [9].

We next evaluate the performance of the LV, MA and LpEMA. Figures 3 and 4 provide the mean prediction error of each flow in the trace for different sizes of the sampling window. In the case of the LpEMA predictor, the results shown are for the best value of α_{max} found for each flow, since its effectiveness depends on this parameter (as Fig. 4 suggests). The results of Figs. 3 and 4 also confirm the previous finding, even when a different trace is used, that the performance of the different predictors does not differ very much in terms of complexity. Nevertheless, these results yield an extra finding, which is, that the errors of all the predictors depends on the granularity of the traffic flows. In fact, in both figures 3 and 4, the prediction errors grow sharply, (roughly) above of the flow for the prefix number 300. Hence, this suggests that the predictors are unable to track very small flows accurately.

Previous finding introduces a practical bound that should be taken into account in manual or adaptive setting of the threshold T. This finding states that, regardless of other factors, it only makes sense to track "popular destinations" if its predictability is effective, i.e., those for which the mean prediction error is bounded. To be more specific, using this error-based prediction criterion with an approximate target maximum error of 100% (i.e., $E(f_i) = \frac{\sum_o^n error_k}{n} \leq 1$, where $E(f_i)$ is the mean prediction error for flow f_i, and $error_k$ is the prediction error in slot of time k), 296 destinations were identified (out of a total of 16150 prefixes), and the sum of their individual traffic represents 99% of the total volume.

Regarding the tuning of LpEMA, it should be noted that the prediction errors for a range of values for α_{max}, flows of different granularity and aggregation, and sizes of the sampling window were examined. Figure 4 shows that the prediction error depends on all these factors. First, it shows that as long as α_{max} increases the prediction error decays significantly, and thus the predictor adapts to the variability of the traffic. When it is above some value for α_{max} the prediction error rises. Second, when the sum of all the traffic, the largest traffic flow, and the smallest traffic flow are compared, it can be observed that the aggregation favors the predictability of the traffic. This is evident for all the predictors. Moreover, the best α_{max} that has been found to track each of the traffic flows is highly variable, which means that a significant effort is required for tuning the LpEMA because there is no common value for all the flows. However, we found that it is more common to use values within the range [1.5, 3.0] or [2.0, 3.5] for bigger traffic flows, in the case of smaller and larger sizes of time intervals respectively.

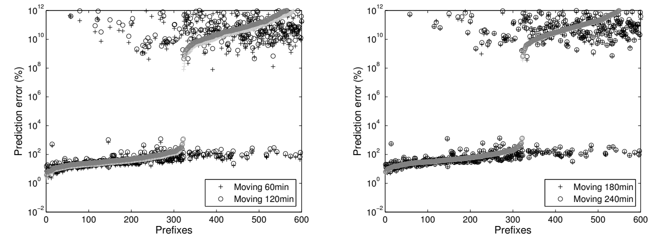

(a) Prediction Errors for Moving Average Predictor.

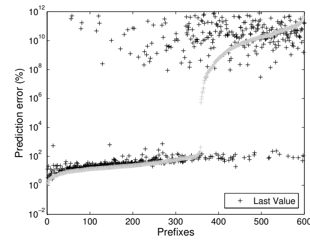

(b) Prediction Error for Last Value Predictor.

Figure 3. Analysis of the Prediction Errors for Moving average and Last value predictors.

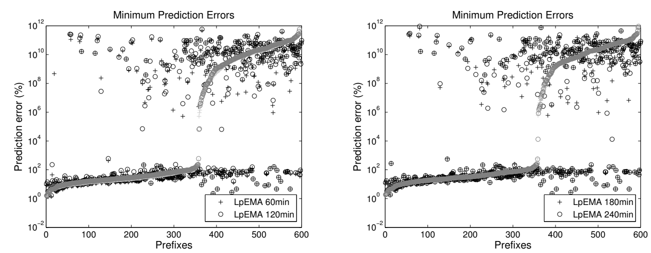

(a) Minimum Prediction Errors.

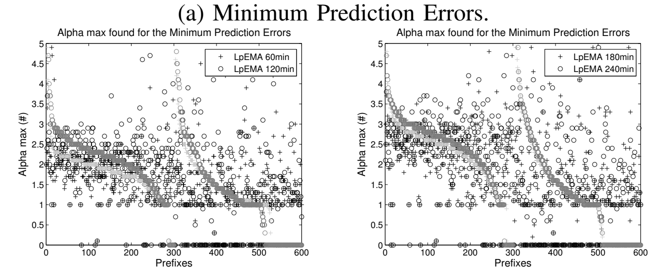

(b) Best Alpha max values found.

Figure 4. Analysis of the Prediction Errors for the LpEMA predictor.

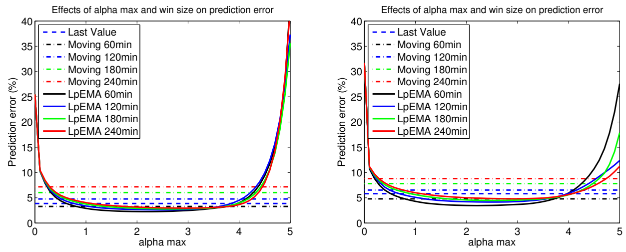

(a) top 50 (left) and top 5 (right) flows.

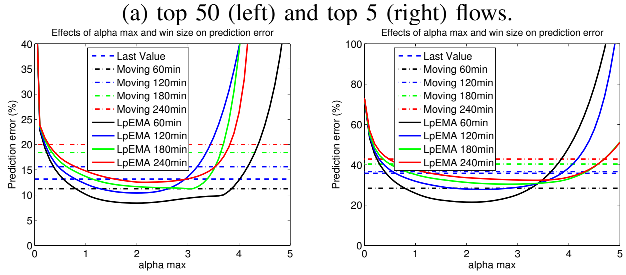

(b) biggest (left) and smallest (right) flows.

Figure 5. Comparison of the prediction errors for different flow sizes and aggregation levels.

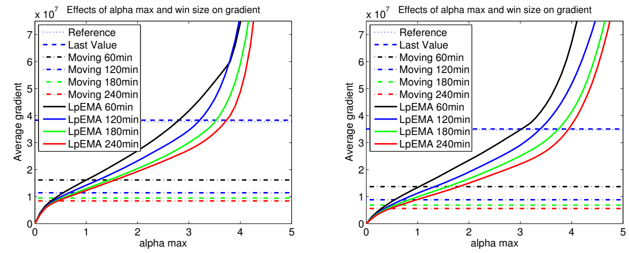

(a) top 50 (left) and top 5 (right) flows.

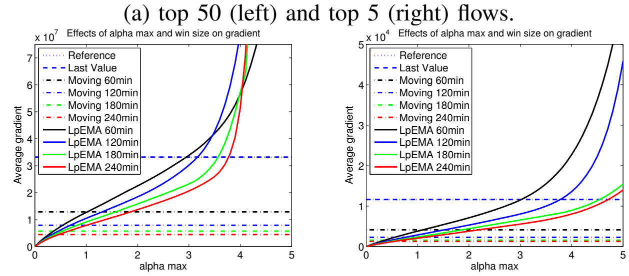

(b) biggest (left) and smallest (right) flows.

Figure 6. Comparison of the average gradients of the samples for different flow sizes and aggregation levels.

One final remark is that the choice of predictors depends on a trade-off between accuracy vs complexity. In [18], the authors, together with their references, argue that the use of complex predictors, such as LpEMA, may not bring performance benefits because their "mean error is always larger than the one of the last value". Figures 5 and 6 provide a counterexample. Our results show that when making proper choices of α_{max}, LpEMA performs always better than LV, and no worse than MA, with regard to mean error. Moreover, the average gradient for LpEMA is always lower than that of LV. In short, the complex predictors may not lead to any benefits in terms of mean error, *but* despite this, may be able to ensure routing stability, while following also the traffic dynamics accurately.

IV. Conclusion

This work has highlighted the importance of the consistency of traffic demands with the Zipf's law for the whole interdomain traffic engineering process integrated in network management systems. This implies that it is enough to take account of only a small fraction of the total number of destinations to control the routing of the majority of the traffic. However, there has been a lack of any simple and pragmatic method for selecting popular destinations.

In view of this, we drew up a practical criterion for the selection of popular destinations. The proposed criterion is based on the definition of a target bound for traffic volumes an relies on information about the behavior of the errors of the traffic predictors. Our results showed that by applying this criterion, we were able to reduce the number of target prefixes to 2% of the total number, while ensuring routing stability due to the predictability of the traffic headed these prefixes.

References

[1] N. Feamster, J. Borkenhagen, and J. Rexford. Guidelines for interdomain traffic engineering. *SIGCOMM Comput. Commun. Rev.*, 2003.

[2] M. Amin, K. Ho, M. Howarth, and G. Pavlou. An integrated network management framework for inter-domain outbound traffic engineering. In *Proceedings of the MMNS'06*, pages 208–222, 2006.

[3] Y. Rekhter, T. Li, and S. Hares. A Border Gateway Protocol 4 (BGP-4). RFC 4271 (Draft Standard), January 2006.

[4] M. Caesar and J. Rexford. BGP routing policies in ISP networks. *IEEE Network*, 19(6), Nov/Dec. 2005.

[5] A. Dhamdhere and C. Dovrolis. ISP and Egress Path Selection for Multihomed Networks. In *Proceedings of the IEEE INFOCOM 2006*, Barcelona, Spain, November/December 2006.

[6] T.C. Bressoud, R. Rastogi, and M.A. Smith. Optimal configuration for BGP route selection. In *Proceedings of the IEEE INFOCOM 2003*, San Francisco, CA, USA, March/April 2003.

[7] S. Uhlig. A multiple-objectives evolutionary perspective to interdomain traffic engineering. *International Journal of Computational Intelligence and Applications*, 2005.

[8] K. Park and W. Willinger. *Self-similar network traffic and performance evaluation*. Wiley-Interscience, 2000.

[9] Lada A. Adamic and Bernardo A. Huberman. Zipf's law and the internet. *Glottometrics*, 3:143–150, 2002.

[10] S. Bhattacharyya, C. Diot, and J. Jetcheva. Pop-level and access-link-level traffic dynamics in a tier-1 POP. In *Proceedings of the 1st ACM SIGCOMM Workshop on Internet Measurement*, 2001.

[11] W. Fang and L. Peterson. Inter-AS traffic patterns and their implications. In *Proceedings of IEEE GLOBECOM '99*, Rio de Janeiro, Brazil, 1999.

[12] A. Feldmann et al. Deriving traffic demands for operational IP networks: methodology and experience. *IEEE/ACM Trans. Netw.*, 9(3), 2001.

[13] S. Uhlig et al. Providing public intradomain traffic matrices to the research community. *SIGCOMM Comput. Commun. Rev.*, 36(1), 2006.

[14] *Cisco IOS Netflow (www.cisco.com/web/go/netflow)*. Accessed 6/2011.

[15] L. Burgstahler and M. Neubauer. New modifications of the exponential moving average algorithm for bandwidth estimation. In *Proceedings of the 15th ITC Specialist Seminar*, Wurzburg, Germany, 2002.

[16] George Edward Pelham Box and Gwilym Jenkins. *Time Series Analysis, Forecasting and Control*. Holden-Day, Incorporated, 1990.

[17] P. Cortez, M. Rio, M. Rocha, and P. Sousa. Internet Traffic Forecasting using Neural Networks. pages 2635 –2642, 2006.

[18] S. Uhlig and O. Bonaventure. Designing BGP-based outbound traffic engineering techniques for stub ASes. *SIGCOMM Comput. Commun. Rev.*, 34(5), 2004.

[19] Qi He, C. Dovrolis, and M. Ammar. On the predictability of large transfer TCP throughput. *Comput. Netw.*, 51(14), 2007.

[20] C.R. Shalizi A. Clauset and M.E.J. Newman. Power-law distributions in empirical data. *SIAM Review, available in http://www.santafe.edu/ aaronc/powerlaws/*, 2009.