# Improving Available Bandwidth Estimation using Averaging Filtering Techniques

Emanuele Goldoni
Dept. of Electronics
University of Pavia, Italy
emanuele.goldoni@unipv.it

Giuseppe Rossi
Dept. of Computer and Systems Science
University of Pavia, Italy
giuseppe.rossi@unipv.it

## ABSTRACT

End-to-end available bandwidth estimation is very important for bandwidth dependent applications, quality-of-service verification and traffic engineering. Although several techniques and tools have been developed in the past, producing precise estimations in real-time still remains challenging, mainly due to the lack of control on the measured network path and environmental conditions, such as noise and short term traffic burstiness. This paper presents VHF (Vertical Horizontal Filter), a modified Exponentially Weighted Moving Average (EWMA) technique borrowed from the financial world; our aim is to produce precise estimates of the available bandwidth, lowering the impact of noisy observations. The accuracy of this solution is compared to other EWMA techniques, showing that VHF has good performance and stability. Simulation results also show that VHF behavior is quite predictable and there is no need to fine tune the filter every time network conditions change.

### Categories and Subject Descriptors

C.2.3 [**Network Operations**]: Network Monitoring

### General Terms

Measurement, Performance, Algorithm

### Keywords

Network measurement, Available bandwidth, Filtering

## 1. INTRODUCTION

The available bandwidth of a network path is a crucial metric in quality-of-service management, traffic engineering in peer-to-peer and overlay networks, congestion control, streaming and VoIP applications. In principle, it would be possible to obtain estimates of the available bandwidth from intermediate routers along the path — this is not feasible in practice, mostly due to technical and security reasons.

Active probing techniques inject end-to-end test traffic into the network and estimate the available bandwidth by observing cross-traffic effects on delays that are experienced by the probe packets. Several active probing tools have been developed and studied in recent years, such as IGI [4], Abing [11], Spruce [14], Pathload [5], TOPP [10, 6], pathChirp [12], FEAT [16] and BART [2]. They differ mainly in the structure or pattern of probe streams and in the methods used to estimate available bandwidth from the observed delays. Many of these tools have been tested in simulation or verified over some Internet paths. However, as noted in [19], there is still great uncertainty in their accuracy and reliability.

The rest of the paper is organized as follows. First, Section 2 briefly illustrates the related work on current filtering techniques. Filters are often adopted by tools, regardless the probing algorithm used, to increase the accuracy of their estimation. This section summarizes the major filtering techniques used in this field, focusing on Vertical Horizontal Filter (VHF), a new filter we propose. Next, evaluation of the proposed filter is presented in Section 3, including simulation results obtained comparing our solution to other similar techniques. Finally, in Section 4 conclusions are drown and future works outlined.

## 2. FILTERING TECHNIQUES

The network model used in many estimation tools is simple and has many shortcomings — e.g. algorithms often ignore that not all routers apply first-in-first-out queues. Moreover, fluid cross-traffic is assumed, disregarding the discrete nature of packets. Other fundamental difficulties with most of the existing tools rise from a number of issues on both end-hosts and network paths: system timing, end-host throughput and end-to-end pathologies [19]. All these factors could introduce a considerable amount of noise in the individual network observations. Figure 1 shows the output of a measurement using pathChirp over a completely empty link: even if no cross-traffic is present, estimation's accuracy is affected by system timing and hardware errors. It should be noted that similar results have been obtained in the past [3] using other well-known tools like Pathload, IGI and Spruce.

The issue of converting noisy observations into an

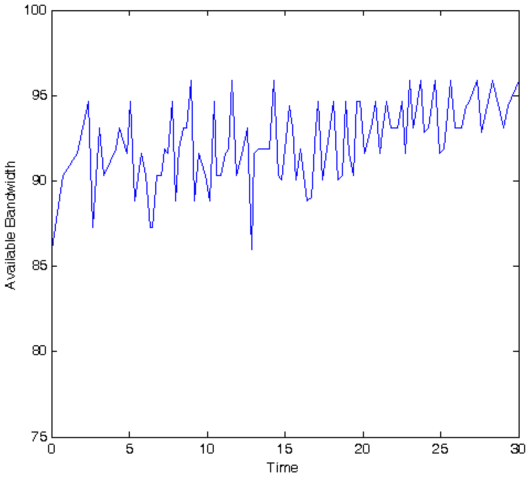

**Figure 1: PathChirp measurement with no cross-traffic.**

estimate is a common filtering problem. The idea of using such a solution in this context is based on the predictability and long-term stability of the Internet. Typically, the available bandwidth of an Internet path shows strong correlation and a certain degree of stability over intervals that span from several minutes to a few hours [18] [16]. Given a new observation, an effective filtering technique should produce a new estimate of the available bandwidth combining both the most recent observation and the old values.

## 2.1 Existing Filters

PathChirp smooths estimations using a moving average filter. This technique computes the current estimation averaging the last $M$ observations, where $M$ is the so-called *window size*. The problem is to find the right size $M$ for the dataset. If the window is too wide, real changes are leveled out and the filter reacts slowly. On the other hand, if $M$ is too small, noisy peaks are not suppressed and the outcome of the filter tends to follow the measured values.

The Exponentially Weighted Moving Average (EWMA) filter uses observed values $O_k$ and outputs a new estimation $E_i$ calculated as follows:

$$E_i = \alpha E_{i-1} + (1-\alpha)O_i. \tag{1}$$

This filter is used by two available bandwidth estimation tools: Abing [11] and Yaz [13]. The authors of the first tool originally set the value of $\alpha$ to 0.9, providing a fairly heavy smoothing. However, the latest version of the software changed the smoothing factor to $\alpha = 0.75$ [9]. Yaz, on the other hand, should use an Exponentially Weighted Moving Average filter with a less strong $\alpha$ ($\alpha = 0.3$), but the adoption of such a technique is only suggested and it has not been actually implemented in the tool.

The difficulty with EWMA filters lies in the choice of the exponential weight $\alpha$. With a large $\alpha$ the old estimates are given more importance and the filter is slow but stable; agility is instead achieved by keeping $\alpha$ small. Neither stability nor agility are desirable all the time. Ideally, the filter should be adaptive, setting the value of $\alpha$ according to the current circumstances. Equation (1) should then take the form:

$$E_i = \alpha_i E_{i-1} + (1-\alpha_i)O_i. \tag{2}$$

The basic idea is that sharp and non-persistent changes can at first be treated as noise using lower weights $\alpha_i$. However, if the change persists, the filter should quickly converge to the new value. *Lowpass EMA* [1], *Stability* [8] and *Error Based* Filters [8] are three existing techniques designed around this philosophy. Although they have been proposed a couple of years ago, to our best knowledge none of them has actively been employed in an available bandwidth estimation tool.

Averaging filters are not the only one used inside available bandwidth estimation tools, although quite common and simple. Kalman filter and Wavelets, sometimes combined with window-based averaging, have been successfully adopted to attenuate noise and local random fluctuations [2] [15] [7]. However, a detailed analysis of these different techniques is outside the scope of this paper

## 2.2 The Vertical Horizontal Filter

We propose Vertical Horizontal Filter (VHF), a new filtering technique borrowed from the financial world [17] but never used for bandwidth estimation. VHF is a modified EWMA filter based on the same principles of the three above mentioned filters. The dynamically exponential weight $\alpha_i$ in (2) is computed as:

$$\alpha_i = \beta \frac{\Delta_{max}}{\sum_{t=i-M}^{i} |O_t - O_{t-1}|} \tag{3}$$

where $\Delta_{max}$ is the gap between the maximum and the minimum values in the $M$ most recent observations. In our model we set $\beta$ as $\frac{1}{3}$ and the window size $M = 10$. These parameters were chosen empirically to minimize estimation error under 10 random simulations of the network scenario presented in Section 3. The weight is statically set to 0.5 if the observed available bandwidth remains constant for at least $M$ consecutive observations — otherwise (3) would create a division by zero.

# 3. SIMULATION RESULTS

In this section we use simulations to compare the Vertical Horizontal Filter with other linear filtering techniques — all applied to the same set of observations.

Real world measurements cannot be reproduced as many times as needed. Moreover, it is often not possible

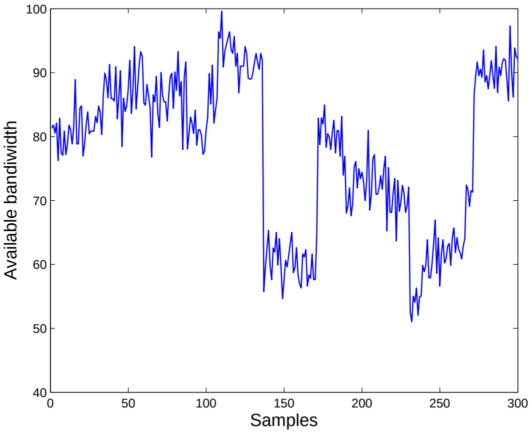

**Figure 2: A typical generated trace file.**

to validate the estimation results due to lack of knowledge of the available bandwidth on the whole network path observed. For these reasons, simulation, thought far from being realistic, is probably the simplest way to evaluate consistently different de-noising algorithms. We have performed a great number of numerical simulations using MATLAB, in order to compare the performance of various smoothing linear filters over the same set of noisy data.

The Vertical Horizontal Filter we propose has been compared to the moving average filter used in pathChirp, a narrower moving average, a standard EWMA and the three modified EWMA filters mentioned in Section 2. In the same Section we have discussed the default values used for the configuration of the VHF filter. As far as the other dynamic EWMA filters are concerned, we have adopted the optimized parameters suggested by the original authors in [8] and [1].

In order to perform our simulation, we have generated 10,000 different available bandwidth traces, each made up of 300 samples; every sequence simulates a variable bit-rate source. We randomly chose the initial transmission rate between 50 Mbps and 95 Mbps and changed it after a random number of samples. We have then added a background source of Gaussian noise and a second source of non-Gaussian noise, with the latter randomly generating spikes in order to simulate traffic bursts. An example of a generated traffic trace is shown in Figure 2.

Since we have both the original traffic traces and the noisy ones, it is possible to calculate how much the filters' estimations differ from the effective available bandwidth value. Our performance metric of choice is the Mean Square Error (MSE) of the estimate for a single trace. An average MSE is obtained averaging all the 10000 errors; the returned value is finally normalized by the average Mean Square Error of the raw, non filtered observations.

Figure 3 demonstrates that most of the moving average filters do *not* increase accuracy. Moreover, this kind of filtering introduces even more error in many cases — a burst, if not effectively identified, can have a negative impact on estimations that follows in time, even if they are not directly affected by the noise.

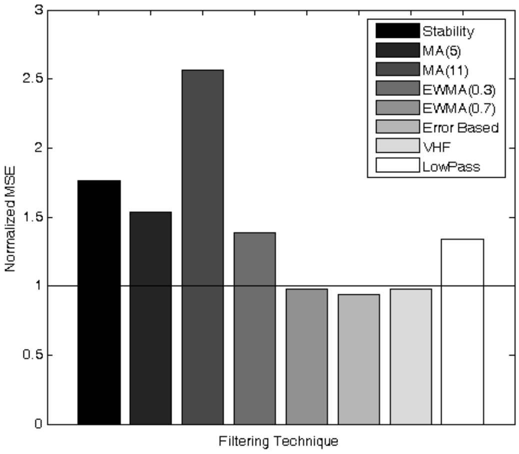

**Figure 3: Comparison between different linear filtering techniques.**

We have repeated the simulations studying the three linear filters that outperformed the other methods proposed. Considering only the Vertical Horizontal Filter and the Error Based Filter, we have studied the relationship between the average Mean Square Error and the number of most recent observations considered. As far as the standard EWMA filter is concerned, we have studied the MSE versus the value of the weight $\alpha$.

Figure 4 shows that the performance of Virtual Horizontal Filter is not significantly affected by the initial setup of the observed window size. On the other hand, Error Based filter performance depends also on the window size parameter and EWMA strongly depends on the value of the weight $\alpha$. Compared to the VHF filter, this could lead to a more complex tuning of the Error Based and EWMA filters as well as network conditions or traffic distribution change.

## 4. CONCLUSIONS AND FUTURE WORK

We presented VHF, a novel linear filter technique that can increase estimations accuracy.

Simulation results' revealed that linear filtering techniques can effectively reduce the impact of noise on the estimation of the available bandwidth. Compared to other similar filtering methods, the VHF filter we propose leads to better results in many cases and shows greater stability. Our results indicate also that there is

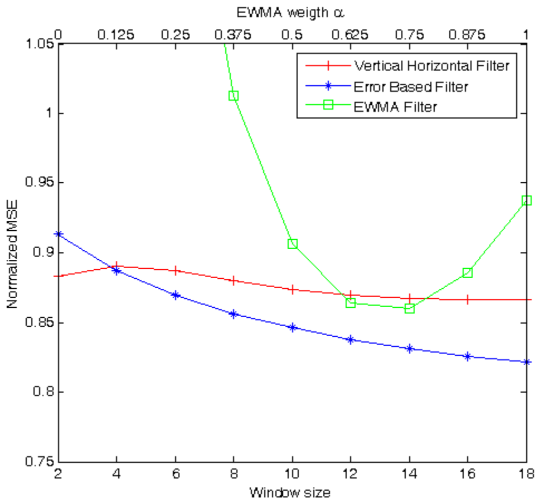

**Figure 4: Normalized MSE under varying window sizes or EWMA weights.**

no need to fine tune our filter every time some network conditions change.

The proposed setup of the VHF filter has been determined empirically: rigorous evaluation would be required to identify, if exists, the optimal configuration of the filter. Moreover, although the outcome of our simulations indicates that the proposed filter is stable, we plan to demonstrate that it is asymptotically stable.

Finally, some preliminary tests have been already conducted in a lab testbed setting. We will test intensively the performance of our solution in the testbed as well as over a limited number of Internet paths, studying the impact of different sources of cross-traffic and of different measurement timescales.

## 5. ACKNOWLEDGEMENTS

The authors would like to thank Dr. Davide Cavalca for giving the paper a critical reading and for providing us several helpful comments.

## 6. REFERENCES

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

[2] S. Ekelin et al. Real-time measurement of end-to-end available bandwidth using kalman filtering. In *NOMS Symposium 2006, 10th IEEE/IFIP*, 2006.

[3] C. D. Guerrero and M. A. Labrador. Experimental and analytical evaluation of available bandwidth estimation tools. In *31st IEEE Conf. on Local Computer Networks*, 2006.

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

[5] M. Jain and C. Dovrolis. Pathload: A measurement tool for end-to-end available bandwidth. In *Proc. of 3rd PAM Workshop*, Fort Collins, 2002.

[6] A. Johnsson at al. Diettopp: A first implementation and evaluation of a simplified bandwidth measurement method. In *2nd Swedish National Computer Networking Workshop*, Karlstad, 2004.

[7] S.-R. Kang and D. Loguinov. IMR-Pathload: Robust available bandwidth estimation under end-host interrupt delay. In *Lecture Notes in Computer Science*, volume 4979/2008, Springer Berlin, 2008.

[8] M. Kim and B. Noble. Mobile network estimation. In *Proc. of MobiCom '01*, New York, 2001.

[9] C. Logg et al. Experiences in traceroute and available bandwidth change analysis. In *Proc. of ACM SIGCOMM Workshop on Network troubleshooting*, New York, 2004.

[10] B. Melander et al. A new end-to-end probing and analysis method for estimating bandwidth bottlenecks. In *Proc. of 43th IEEE GLOBECOM Conf.*, San Francisco, 2000.

[11] J. Navratil and R. L. Cottrell. Abwe: A practical approach to available bandwidth. In *Proc. of 4th PAM Workshop*, San Diego, April 2003.

[12] V. Ribeiroet al. pathchirp: Efficient available bandwidth estimation for network paths. In *Proc. of 4th PAM Workshop*, San Diego, April 2003.

[13] J. Sommers et al. A proposed framework for calibration of available bandwidth estimation tools. In *Proc. of 11th IEEE Symposium on Computer and Communications*, 2006.

[14] J. Strauss et al. A measurement study of available bandwidth estimation tools. In *Proc. of 3rd Conf. on Internet Measurement*, New York, 2003.

[15] G. Urvoy-Keller et al. Operational comparison of available bandwidth estimation tools. *SIGCOMM Comput. Commun. Rev.*, 38(1), 2008.

[16] Q. Wang and L. Cheng. FEAT: Improving accuracy in end-to-end available bandwidth measurement. In *Proc. of 49th IEEE GLOBECOM Conf.*, Nov. 2006.

[17] A. White. The vertical horizontal filter. *Futures Magazine*, Volume XX, No. 10:1–10, 1991.

[18] Y. Zhang and N. Duffield. On the constancy of internet path properties. In *Proc. of 1st Workshop on Internet Measurement*, New York, 2001.

[19] H. Zhou et al. Difficulties in estimating available bandwidth. In *Proc. of IEEE ICC'06 Conf.*, 2006.