

Estimating bandwidth requirements using flow-level measurements

Ricardo de O. Schmidt, Aiko Pras

University of Twente
Design and Analysis of Communication Systems (DACS)
Enschede, The Netherlands
`{r.schmidt,a.pras}@utwente.nl`

Abstract. Bandwidth provisioning is used to dimension links so that a desired performance is met within a network (e.g. QoS metrics). Practical approaches for the estimation of bandwidth needs rely on rough traffic measurements overlooking traffic fluctuations that may have a negative impact on the network performance. On the other hand, effective theoretical solutions for estimating bandwidth needs require traffic measurements at very small timescales, which limit their applicability due to, for example, transmission rates. This paper describes our research goal of proposing a reliable and scalable solution for estimating bandwidth requirements by means of flow-level traffic measurements, as well as our proposed approach to achieve it. This research, currently in its initial phase, is planned to result in a Ph.D. thesis at the end of four years.

1 Introduction and motivation

Link dimensioning plays an important role in network management, for example to support QoS provisioning. Usually, network operators make manual provisioning decisions based on empirical rules and on rough traffic measurements, like by using the Simple Network Management Protocol (SNMP) [1], and the Multi Router Traffic Grapher (MRTG) [2] for graphics, in timescales of 5 minutes. However, as proved in [3], traffic measurements in large timescales (e.g. interval of minutes) completely overlook traffic peaks that are detected in smaller timescales (e.g. intervals of milliseconds). This leads to imprecise estimations of bandwidth needs and, consequently, it results in non-optimal network performance. Both underdimensioning and overdimensioning of links can impact negatively in the network performance. The former can result in traffic bottlenecks caused by undetected traffic peaks, and the latter in waste of link resources when allocating more bandwidth than necessary.

Some of the solutions proposed so far for network dimensioning are application specific and, consequently, have reduced applicability scope. Solutions like [4] and [5] are focused on delay sensitive applications (e.g., VoIP) in backbone links, where the delay is resulted from queuing. By means of flow-level measurements, [6] and [7] focus on, respectively, dimensioning of TCP-connections throughput rates and congestion control.

A more general approach, i.e. applicable not only for backbone traffic, is proposed in [8] and [9]. One of the main contributions of this work is the dimensioning formula (1). It is used for calculating the minimal bandwidth requirements C for a given link, in a time scale of length T and with an acceptable error rate of ε , by adding to the average traffic rate μ a safety margin, which strongly depends on the traffic variance $\nu(T)$. The assumption of this formula is that the measured traffic at timescale T is Gaussian.

$$C(T, \varepsilon) = \mu + \frac{1}{T} \sqrt{(-2 \log \varepsilon) \cdot \nu(T)} \quad (1)$$

This formula was extensively validated and proved to be a good solution for the problem of link dimensioning. However, the drawback of this approach is that it requires packet-level traffic measurements, which may drastically reduce its applicability in high-speed networks (e.g., Gigabit Ethernet). Considering traffic loads currently observed in the Internet, packet-level measurements could often suffer from scalability issues. In addition, it may demand very specific, powerful and expensive hardware/software for traffic monitoring.

Given the increasing employment of flow-level measurements on network management and operations, e.g. Cisco IOS routers with embedded NetFlow [10], we believe it to be a promising approach for our research problem. Considering current transmission rates, the use of flow-level traffic measurement tools (e.g., NetFlow and IPFIX [11]) would allow us to come up with a scalable solution when working with aggregated traffic information. Therefore, this paper describes our research idea towards the proposal of a solution for reliably estimating bandwidth requirements without recurring to packet-level traffic measurements.

The remaining of this paper is organized as follows. Section 2 presents our goal and research questions, as well as the proposed approach. The planned process to validate our findings is described in Section 3. Finally, Section 4 concludes the paper.

2 Goal, research questions and approach

The goal of our research is *to determine how to estimate bandwidth requirements by means of flow-level traffic measurements*. In order to pursue this goal, we have defined the following research questions as the base of our research:

1. *What is the state of the art in flow-based bandwidth provisioning?*
2. *Which relevant data extracted from flow records can be used to estimate bandwidth needs?*
3. *Which statistical models describe the traffic measured at the flow-level?*

Research question 1 aims to raise information about main contributions and trends in the area of our research. To do so, we plan an exhaustive literature

review. Interviews with network operators can also be considered as a possible way for obtaining information from real-world situations.

With research question 2, we want to identify what information (e.g., metrics) is relevant to our problem and how to extract it from the flow records. For instance, we aim to find out how to determine average traffic rate and variance from flow measurements in order to further apply these values in the dimensioning formula (1), as presented in section 1. It is important to remark that, as presented by [8], determining traffic variance from packet-level measurements (with smaller timescales) is not a complex task, due to the high level of information granularity, which allows for a high accuracy. However, the same task can be considered as a challenge in flow-level measurements, where traffic information is summarized into flow records and measurements occur at coarser timescales. In order to answer research question 2, we will use the knowledge acquired from research question 1 and further analysis on flow-level measurements of real-traffic.

In research question 3, we aim at identifying which traffic models fit better the traffic measured at the flow-level. More specifically, we want to investigate whether the measured traffic is Gaussian or not, in order to determine the suitability of the dimensioning formula from [8] for estimating bandwidth needs from flow-level traffic measurements. Once the traffic model(s) is identified, it will allow us to parameterize metrics such as average traffic rates and peaks (see [12]). We plan to use real-world network traces in order to identify such models.

3 Validation

In order to validate our final proposal, we intend to evaluate its correctness by analyzing real-world traffic, from different sources like GÉANT [13] and SURFnet [14]. We plan to compare the results from our solution with results obtained from the dimensioning formula in [8] with packet-level measurements. In addition, experimentations with real equipment and real-world traffic measurements will be done in the laboratory.

4 Final considerations

This paper discusses the possibility of estimating network bandwidth requirements by means of flow-level traffic measurements. We have stated our motivation and planned means for achieving our goal. The long-term research goal, as described in this paper, is to be achieved in the period of 4 years, as a Ph.D. research. The research reported in this paper is supported by the SURFnet [14] and the FP7 ICT UniverSelf [15] (#257513).

References

1. Schonwalder, J.: Simple Network Management Protocol (SNMP) Context EngineID Discovery. RFC 5343, 2008.

2. MRTG: Multi Router Traffic Grapher. Available at: <http://oss.oetiker.ch/mrtg/>. Accessed in: Feb. 2011.
3. Meent, R. van de, Pras, A., Mandjes, M., Berg, H. van den, Nieuwenhuis, L.: Traffic Measurements for Link Dimensioning, A Case of Study. In proceedings of the 14th IFIP/IEEE Workshop on Distributed Systems: Operations and Management (DSOM), pp. 106–117, 2003.
4. Fraleigh, C.: Provisioning Internet Backbone Networks to Support Latency Sensitive Applications. PhD thesis, Stanford University, 2002.
5. Papagiannaki, K.: Provisioning IP Backbone Networks Based on Measurements. PhD thesis, University of London, 2003.
6. Bonald, T., Olivier, P., Roberts, J.: Dimensioning high speed IP access networks. In proceedings of the 8th International Teletraffic Congress (ITC), pp. 241–251, 2003.
7. Barakat, C., Thiran, P., Iannaccone, G., Diot, C., Owezarki, P.: Modeling Internet backbone traffic at the flow level. *IEEE Transactions on Signal Processing*, 51(8), 2003.
8. Meent, R. van de: Network Link Dimensioning - a measurement & modeling based approach. PhD thesis, University of Twente, 2006.
9. Pras, A., Nieuwenhuis, L., Meent, R. van de, Mandjes, M.: Dimensioning Network Links: A New Look at Equivalent Bandwidth. *IEEE Network*, 23(2), pp. 5–10, 2009.
10. Cisco Systems Inc.: Cisco IOS Flexible NetFlow Configuration Guide. Release 12.4T, 2008. Available at: <http://www.cisco.com/>. Accessed in: Feb. 2011.
11. Quittek, J., Zseby, T., Claise, B., Zander, S.: Requirements for IP Flow Information Export (IPFIX). RFC 3917, 2004.
12. Meent, R. van de, Pras, A., Mandjes, M. R. H., van den Berg, J. L., Roijers, F., Nieuwenhuis, L. J. M., Venemans, P. H. A.: Burstiness predictions based on rough network traffic measurements. In proceedings of the 19th World Telecommunications Congress (WTC/ISS), 2004.
13. GÉANT: Géant. Available at: <http://www.geant.net/>. Accessed in: Feb. 2011.
14. SURFnet: Surfnet. Available at: <http://www.surfnet.nl/>. Accessed in: Feb. 2011.
15. UniverSelf: UniverSelf FP7 project number 257513. Available at: <http://www.univerself-project.eu/>. Accessed in: Feb. 2011.