EuroNGI Workshop on QoS and Traffic Control, Paris, 7-9 December 2005

RISK REDUCTION IN THE HOSE MODEL FOR VPN DESIGN

Maurizio Naldi
Università di Roma "Tor Vergata"
Dip. di Informatica Sistemi Produzione

EXTENDED ABSTRACT

Network dimensioning is traditionally approached through the consideration of traffic matrices: the operator is assumed to know the traffic intensity for each of the Origin-Destination couples in its network. Traffic matrix measurements have therefore a central role in network planning and significant efforts have been spent, first in the telephone network and now in IP networks [1] [2], to measure the O-D traffic as accurately as possible, over a time period as extended as possible and with a suitable time granularity.

However, in recent years, significant efforts have been spent in investigating more relaxed dimensioning approaches as to the requirement on traffic matrix knowledge. This has been done on several grounds, namely that: a) the accuracy exhibited by the traffic matrix (often obtained by indirect or non continuous measurements) is not very high; b) the variability of the traffic matrix is very high; c) the traffic matrix is not available.

An approach that has been proposed and has met the favour of network designers is the so-called hose model, where the designer is assumed to know the overall incoming and the overall outgoing traffic for each node, rather than all its constituent streams (i.e. the contributions of every other ingress node for the incoming traffic and of every other egress node for the outgoing traffic). This model, which can be seen a special case of the polyhedral model [7], has been particularly welcome for the design of Virtual Private Networks, since the VPN owner typically knows its overall traffic but doesn't possess a deep knowledge of the associated origins or destinations [3]. However, this provides the hose model with a rationale just for the cases described by point c) of the above recalled arguments.

A benefit of the hose model that has instead be neglected is that considering traffic at a higher aggregation level than that associated to O-D streams, it allows for a statistical multiplexing gain and risk reduction in the evaluation of the reference traffic to be used for planning purposes.

In fact, the traffic values to be input to the dimensioning procedure is usually a collective reference value to be computed based on a number of traffic measurements for the same entity (e.g. the same transmission link or the same O-D couple) over an extended period.

Different ways have been proposed to derive the reference value from the fine-grain measurements (e.g. collected every 5 minutes), typically based on a combination of averaging and order statistic extraction operations [4] [5]. The resulting value must account for the very most part of traffic intensities but avoid being set too high because of very rare extreme values.

The computation of the traffic reference value can be seen as a risk-protection measure, since the dimensioning procedure guarantees the desired QoS level when the actual traffic is lower or equal to the traffic reference value. The traffic reference value is therefore set so that the probability of exceeding it, i.e. the probability of failing the QoS objectives, agrees with the operator's policy. As in all risk-related settings larger variability levels are associated with larger risks: any measure which tries to reduce the variability associated to the quantities of interest is therefore a risk-reducing measure. As long as it helps in reducing the traffic variability the hose model can therefore act as a risk-reducing approach, an advantage which can somewhat offset the price to be paid for the lack of O-D traffic knowledge.

In order to measure the risk-exposure associated to the single O-D streams vs. the overall incoming or outgoing traffic we have considered two measures of variability and have analysed the behaviour of two real networks, as embodied by their associated traffic matrices gathered for an extended period of time, kindly supplied by Vaton [6], which concern two small networks, respectively of 6 and 13 nodes.

Let's consider the risk associated to the variability of traffic. The more the traffic is variable the more it calls for a conservative dimensioning. This dimension of risk may therefore be captured by the coefficient of variation of the random variable that measures the traffic intensity. Indicating by T_{ij} the traffic for the O-D couple made by the sending node *i* and the receiving node *j*, the risk associated to this traffic relationship is the ratio

$$CV_{ij} = \frac{\sqrt{Var\left[T_{ij}\right]}}{E\left[T_{ij}\right]}$$

When applying the hose model we are not concerned anymore with the detailed streams pertaining to each O-

---

D couple, but we rather consider the whole originated (or terminated) traffic. For each sending node we are therefore interested in the random variable

$$T_i = \sum_{j \neq i} T_{ij}$$

The resulting risk associated to the traffic originated by this node is now given by the coefficient of variation of this random variable, which can be related to the constituents of the risk associated to the single destination streams:

$$CV_i = \frac{\sqrt{Var[T_i]}}{E[T_i]} = \frac{\sqrt{\sum_{j \neq i} Var[T_{ij}] + \sum_{j \neq i} \sum_{k \neq i,j} Cov(T_{ij}, T_{ik})}}{\sum_{j \neq i} E[T_{ij}]}$$

Depending on the covariance structure, this risk can actually be lower than the risks associated to the single destination streams, i.e. the single elements of the traffic matrix: taking into account the aggregates (by row and columns) rather the single elements of the traffic matrix can therefore provide the hose model with some gain. It is to be remarked that the risk reduction doesn't require a negative correlation between the O-D streams; even significant positive correlation may give rise to a variability reduction.

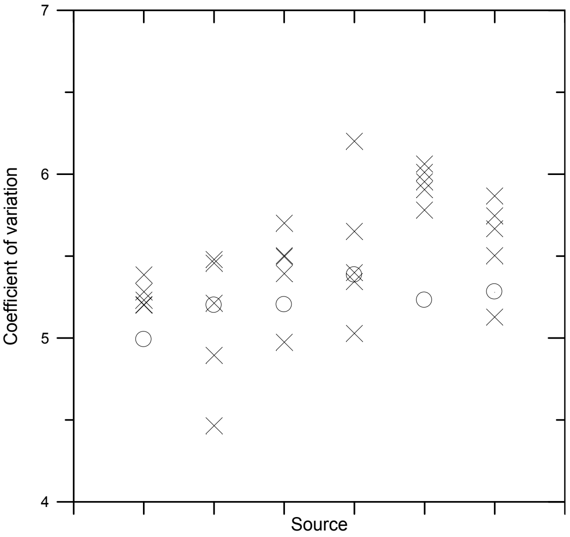

Figure 1 – Coefficient of variation for a sample network

A sample result is provided in Figure 1 where for each of the 6 source nodes the coefficient of variation of the whole source (represented by a circle) and of the O-D streams corresponding to the 5 destinations (represented by a cross) are reported.

As can be seen the sources present a CV lower than the single O-D stream in a large majority of the cases (actually in 24 cases out of 30, i.e. 80%), so that working at the source level allows a strong risk reduction.

Another way of comparing the risk exposure is to examine the behaviour of order statistics, since these are a fundamental components of basically all the computation methods for the traffic reference value. A typical percentile value used for the traffic reference value is the 95% one [5]. We can therefore compare the percentile pertaining to the source node with the sum of those pertaining to the different destination streams, concluding that there is a reduction of the risk exposure if the source node percentile results lowers. The experimental data show that this is indeed the case.

REFERENCES

[1] ITU-T: Recommendation E.501 "Estimation of traffic offered in the network", Geneva, May 1997

[2] A. Medina, N. Taft, K. Salamatian, S. Bhattacharyya, C. Diot: "Traffic Matrix Estimation: Existing Techniques and New Directions", ACM SIGCOMM. Pittsburgh, USA. August 2002

[3] N.G. Duffield, P. Goyal, A.G. Greenberg, P.P. Mishra, K.K. Ramakrishnan , Jacobus E. van der Merwe: "Resource management with hoses: point-to-cloud services for virtual private networks", IEEE/ACM Transactions on Networking, November 2002

[4] ITU-T: Recommendation E.500 "Traffic intensity Measurement Principles", Geneva, 1998

[5] N. Brownlee, R. Fulton: "Hawaihiko and the Third-Quartile Day", IEEE Communications Magazine, August 2000, pp. 162-168

[6] S. Vaton: Personal communication, May 2004

[7] W. Ben-Ameur, H. Kerivin: "Routing of Uncertain Demands", Proceedings of the 33th Annual ACM Symposium on Theory of Computing pages 389-398, 2001.