Scalable accurate consolidation of passively measured statistical data

Silvia Colabrese[1], Dario Rossi[1], Marco Mellia[2]

[1]Telecom ParisTech, Paris, France - {silvia.colabrese,dario.rossi}@enst.fr
[2]Politecnico di Torino, Torino, Italy - marco.mellia@polito.it

Abstract. Passive probes continuously collect a significant amount of traffic volume, and autonomously generate statistics on a large number of metrics. A common statistical output of passive probe is represented by probability mass functions (pmf). The need for consolidation of several pmfs arises in two contexts, namely: (i) whenever a central point collects and aggregates measurement of multiple disjoint vantage points, and (ii) whenever a local measurement processed at a single vantage point needs to be distributed over multiple cores of the same physical probe, in order to cope with growing link capacity. Taking an experimental approach, we study both cases assessing the impact of different consolidation strategies, obtaining general design and tuning guidelines.

1 Introduction

This paper focuses on consolidation of multiple statistics, and especially of distribution quantiles, gathered from passive probes. We briefly mention two completely orthogonal scenarios where this need arises. First, in the case of multiple vantage points, consolidation of data coming from multiple sources yields a more statistically representative population sample. Second, in the case of a single vantage point, it may be necessary to split traffic processing over multiple independent cores to avoid CPU bottlenecks.

These two scenarios appear to be rather different at first sight. In the former case, the number of vantage points can vary between a handful to many different collection points, each of which gathers traffic mixture with likely different characteristics. In the latter case, processing can be split among a handful of cores in a CPU (possibly more for GPU-based architectures), each of which is processing a random sample of the incoming traffic.

Yet, on a closer look these diverse scenarios translate into similar constraints. In the case of multiple vantage points, the first and foremost constraint is represented by the amount of data that is required for the consolidation – transferring the least possible amount of data is hence desirable. In the case of parallel processing at a single vantage point, the constraint is instead represented by processing power – to limit the computational overhead tied to the consolidation process, elaborating the least possible amount of data is hence desirable.

Hence, we argue that a single flexible methodology could fit both purposes, which we address in this paper.

Fig. 1. Synopsis of the experimental workflow

2 Methodology

While from computational complexity or network overhead viewpoints, reducing the amount of data to be processed and transferred would be desirable, this however clearly tradeoffs with accuracy: in this paper, we focus on this tradeoff.

While our aim is to obtain general design and tuning guidelines, our experiments are based on a specific instance of metrics gathered through the Tstat measurement tool, a passive flow-level monitor that we developed over the last years [4]. For each flow, Tstat tracks over 100 metrics (see [2] for more details), that are used to build standard fixed-width histograms. Percentiles of the distribution are then evaluated with linear interpolation, and stored in Round Robin Databases (RRD).

Our methodology is as in Fig. 1. Input blocks (shaded gray) are quantiles vectors q_i gathered from multiple (local or remote) probes, which are processed to gather consolidated quantile vectors as output to the process. We *interpolate* input quantile vector to get a cumulative distribution function (CDF). We consider two interpolation strategies, namely: a *Linear (L)* and a *Monotonic Spline (S)* strategy (in the latter case, we ensure monotonicity using the Piecewise Cubic Hermite Interpolating Polynomial [5]).

These interpolated functions are weighted by the amount of traffic they represent (weights can be computed in terms of flows, packets or bytes), and added to get the total CDF. Finally, as output of our workflow, we obtain the consolidated $\overline{q}_i$ deciles vector from the total CDF with the bisection method. Finally, the consolidated continuous output is compared to the real quantiles $\tilde{q}_i$ of the aggregated distribution, obtained from running Tstat on the aggregated traces. In what follows, we evaluate the accuracy of the overall workflow by assessing the relative error $(\overline{q}_i - \tilde{q}_i)/\tilde{q}_i$.

As ouput, we limitedly consider *deciles* of the distribution. However, our methodology exploits input quantiles to reconstruct the distributions and operate over CDFs, so that in principle input and output sets do not need to be homogeneous. We thus argue that CDF interpolation can benefit of a larger number of samples (i.e., knots in Spline terms), providing a more accurate description for the intermediate consolidation process. As such, we consider two cases: a *Single (S)* case, where *deciles* are both input and output, and a *Double (D)* case where we additionally use intermediate quantiles (i.e., 5th, 15th, 25th to 95th) as input to the process.
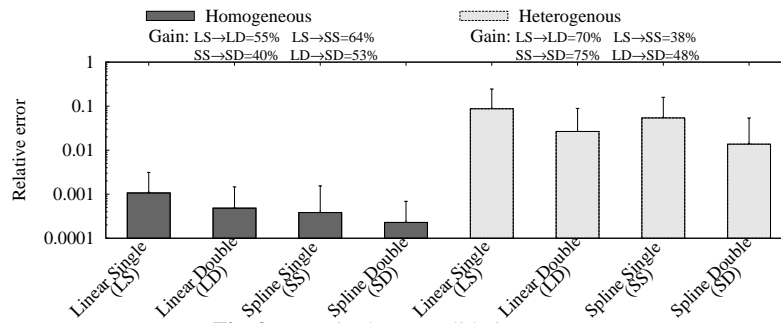
Fig. 2. Error in the consolidation process

3 Experiments

We use several traces, some of which are publicly available. Vantage points pertain to different network environments (e.g., Campus [1, 3] and ISP networks [6]), countries (e.g., EU [3,6] and Australia [1]) and have been collected over a period of over 8 years.

We compactly represent consolidation errors in Fig. 2 (meand and stdev bars over all metrics and quantiles), indicating with *homogeneous* and *heterogeneous* the case of multiple local and distributed vantage points respectively. We further annotate the picture with relative accuracy gain with respect to different consolidation strategies. Shortly, (i) consolidation error is practically negligible for local processes (median error is about 0.1% and maximum 1%), but large for heterogeneous probes (median 1%, maximum 30%, and possibly >100% for naïve strategies); (ii) the use of intermediate quantiles (e.g., 5th, 15th, and so on), is desirable as it significantly improves accuracy (up to 75% in the case of multiple vantage points); (iii) interpolation via Splines is preferable, as it yields to an accuracy gain over Linear interpolation of 40% in our dataset.

Acknowledgement

References

1. Auckland traces. `http://www.wand.net.nz/`.
2. Tstat homepage. `http://tstat.tlc.polito.it`.
3. Unibs traces. `http://www.ing.unibs.it/ntw/tools/traces/`.
4. A. Finamore, M. Mellia, M. Meo, M. Munafo, and D. Rossi. Experiences of internet traffic monitoring with tstat. *IEEE Network*, 2011.
5. F. N. Fritsch and R. E. Carlson. Monotone piecewise cubic interpolation. *SIAM Journal on Numerical Analysis*, 17(2):238–246, 1980.
6. D. Tammaro, S. Valenti, D. Rossi, and A. Pescape. Exploiting packet sampling measurements for traffic characterization and classification. *Wiley IJNM*, pages 451–476, 2012.