Successive c-Optimal Designs : A Scalable Technique to Optimize the Measurements on Large Networks

Guillaume Sagnol
INRIA Saclay & CMAP, École Polytechnique
guillaume.sagnol@inria.fr

Mustapha Bouhtou
Orange Labs
mustapha.bouhtou@orange-ftgroup.com

Stéphane Gaubert
INRIA Saclay & CMAP, École Polytechnique
stephane.gaubert@inria.fr

ABSTRACT

We propose a new approach to optimize the deployment and the sampling rates of network monitoring tools, such as Netflow, on a large IP network. It reduces to solving a stochastic sequence of Second Order Cone Programs. We validate our approach with experiments relying on real data from a commercial network.

Categories and Subject Descriptors: C.2.3 [Network Operations]: Network monitoring

General Terms: Measurement.

1. INTRODUCTION

Internet providers collect traffic measurements for multiple purposes, such as traffic engineering, performance, security or billing. In the present paper, we address the problem of optimizing the measurements for the estimation of the traffic matrix of a large backbone network, which describes the volume of traffic between each two nodes (routers). We believe that this approach might also be of some interest for other purposes, since it indicates which routers or interfaces meet a maximal proportion of the traffic. Moreover, we will see that this approach leads to a nice mathematical formulation and to scalable algorithms.

We consider measurements collected by a network-monitoring tool such as Netflow, a technology of Cisco. Activating Netflow everywhere on the network yields an extensive knowledge of the Origin-Destination (OD) flows. However, it is of great interest to optimize the use of this tool, because it sends heavy records of the sampled packets through the network, which causes an overhead in terms of CPU utilization and bandwidth consumption.

2. THE MODEL AND OUR ALGORITHM

We denote by $\boldsymbol{x}$ the vector of the traffic flow volumes over a given time window. Netflow records contain the IP address of the source and the destination of the sampled packets. This can be used together with the internal routing tables of the network to find to find the "internal destination" of the packets, i.e. the router through which they will exit the network of interest. By contrast, finding the "internal source" of the sampled packets is a challenging issue, and so we assume that Netflow measurements break out the flows traversing a given interface according to their (internal) destination [1]. On the interface k, this results in a observation $\boldsymbol{y}_k$, whose entry d is the sum of the flows traversing k which have destination d (up to a sampling noise):

$$\boldsymbol{y_k} = A_k \boldsymbol{x} + \boldsymbol{\varepsilon_k}. \tag{1}$$

In [1] we show that both the problem of finding an optimal set of locations for Netflow and the problem of setting optimal sampling rates in the monitoring system can be formulated in the form:

$$\max_{\boldsymbol{w}\in\mathcal{W}} \Phi(M(\boldsymbol{w})), \quad \text{where } M(\boldsymbol{w}) := \sum_{k\in[s]} w_k A_k^T A_k, \tag{2}$$

and Φ is an information function from the *optimal experimental design* theory. A typical choice for Φ is the $A-$optimality criterion, $\Phi_A : X \mapsto \operatorname{trace}(X^{-1})$, which has been used for an optimal sampling problem in [3]. For the optimal deployment problem, the variable $\boldsymbol{w}$ is a $0/1-$ vector which indicates on which interface we activate Netflow, and for the optimal sampling problem, $\boldsymbol{w}$ is the vector of sampling rates. The set $\mathcal{W}$ reflects some constraints imposed by the operator for the use of Netflow. In this short paper, we restrain ourselves to the case of a single linear constraint $\boldsymbol{p}^T\boldsymbol{w} \leq b$, but we point out that our model was extended in [1] to handle more realistic, *per-router* constraints, which consist in a set of linear inequalities of the form $R\boldsymbol{w} \leq \boldsymbol{b}$, and to take into accounts the link measurements (SNMP).

While semidefinite programming (SDP) can be used to solve the $A-$optimal design problem [3], it becomes intractable by the state-of-the-art solvers when the network grows in dimension (the typical limit is around $n = 15$ nodes, because the variable of this SDP is a $n^2 \times n^2$ matrix). Instead, we propose to solve a succession of $\boldsymbol{c}-$optimal design problems, which are usually well-suited if we want to estimate a linear combination $\boldsymbol{c}^T\boldsymbol{x}$ of the elements of the traffic matrix. The $\boldsymbol{c}-$optimal design problem is defined through the information function $\Phi_{\boldsymbol{c}} : X \mapsto \boldsymbol{c}^T X^\dagger \boldsymbol{c}$, where $\dagger$ denotes the Moore-Penrose inverse. We next show that this problem can be solved by Second Order Cone Programming (SOCP), which can be done very efficiently by interior point codes.

THEOREM 1. *Let $\boldsymbol{w}^*$ be the optimal solution of the continuous $\boldsymbol{c}-$optimal design problem:*

$$\min_{\boldsymbol{w}\geq 0} \boldsymbol{c}^T M(\boldsymbol{w})^\dagger \boldsymbol{c} \qquad s.t.\ \boldsymbol{p}^T\boldsymbol{w} \leq b, \quad M(\boldsymbol{w}) = \sum_{k\in[s]} w_k A_k^T A_k.$$

And let $\boldsymbol{h}^*, (\boldsymbol{\mu}^*, \boldsymbol{z}^*)$ be a pair of primal and dual solution of the SOCP given in primal and dual form:

$$\max_{\boldsymbol{h}\in\mathbb{R}^m} \boldsymbol{c}^T\boldsymbol{h} \qquad \forall i \in [s], \quad \|A_i\boldsymbol{h}\| \leq \sqrt{\frac{p_i}{b}}$$

$$\min_{\boldsymbol{\mu}>0,(\boldsymbol{z_i})_{i\in[s]}} \sum_{i\in[s]} \mu_i \qquad \sum_{i\in[s]} \sqrt{b/p_i} A_i^T \boldsymbol{z_i} = \boldsymbol{c} \qquad \forall i \in [s], \|\boldsymbol{z_i}\| \leq \mu_i.$$

SIGMETRICS'10, June 14–18, 2010, New York, New York, USA.
ACM 978-1-4503-0038-4/10/06.

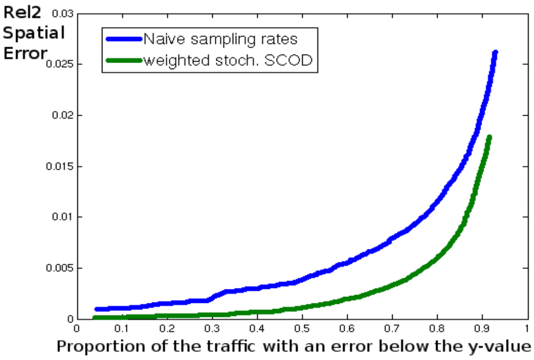

Figure 1: Spatial distribution of the L_2−error for the optimal sampling problem on Opentransit.

Setting $T = \sum_{i \in [s]} \mu_i^*$, *the following relations hold :*

$$\forall i \in [s] \;\; w_i^* = \mu_i^* b/(p_i T), \quad \boldsymbol{c}^T M(\boldsymbol{w}^*)^\dagger \boldsymbol{c} = T^2$$

This theorem is proved in [2]. We propose here to solve the SOCP of Theorem 1 for several vectors $\boldsymbol{c}$, and then to combine the resulting $\boldsymbol{c}$−optimal designs by taking the mean. An interesting choice is to draw the vectors $\boldsymbol{c}$ from a normal distribution $\mathcal{N}(\mathbf{0}, I)$; in this case our algorithm is a Monte-Carlo approximation of the vector

$$\mathbb{E}[\operatorname{argmin}_{\boldsymbol{w} \in \mathcal{W}} \; \boldsymbol{c}^T M(\boldsymbol{w})^\dagger \boldsymbol{c}]. \tag{3}$$

We denote this scheme by SCOD (Successive $\boldsymbol{c}$−optimal designs). Permuting the expectation operator $\mathbb{E}[\cdot]$ with argmin$(\cdot)$ is not permitted in general, but if we do so we end up with the A-optimal design problem, by noticing that $\mathbb{E}[\boldsymbol{c}\boldsymbol{c}^T] = I$:

$$\min_{\boldsymbol{w} \in \mathcal{W}} \operatorname{trace}(\mathbb{E}[\boldsymbol{c}\boldsymbol{c}^T] M(\boldsymbol{w})^{-1}) = \min_{\boldsymbol{w} \in \mathcal{W}} \operatorname{trace} M(\boldsymbol{w})^{-1},$$

We do not claim that our algorithm converges to a A−optimal design, but the latter remark somehow gives sense to the SCOD process. Moreover, it was observed on several examples that the SCOD "converges" rapidly to a design which is close to the A−optimal design (see Table 1). We also define a weighted process (WSCOD), where the vectors $\boldsymbol{c}$ are drawn with respect to the law $\mathcal{N}(\mathbf{0}, \operatorname{diag}(\hat{\boldsymbol{x}}))$, where $\hat{\boldsymbol{x}}$ is a prior estimate of the flow vector $\boldsymbol{x}$.

Design ($\times 10^{-1}$)	SCOD ($N = 10$)	SCOD ($N = 50$)	A-optimal (SDP)
CPU (sec.)	3.72	18.7	492.6
w_1 (Atlanta)	0.559	0.779	0.749
w_3 (Denver)	1.721	1.592	1.510
w_5 (Indiana)	1.458	1.291	1.361
w_7 (LA)	0.556	0.572	0.657
w_{10} (Seattle)	0.000	0.002	0.000

Table 1: Comparison of the A-optimal design and SCOD computed by averaging $N = 10$ and $N = 50$ $\boldsymbol{c}$−optimal designs, for a deployment instance on the Abilene Network with the constraint $\sum_i w_i \leq 1$ (5 out of 11 routers are displayed).

3. VALIDATION

We use real data from the commercial Opentransit network of France Telecom, which consists in 116 nodes, 13456 OD pairs, and 436 links. We simulate the measurements. Our experimental setting is completely described in [1]. We evaluate the deployment of Netflow by the Relative L_2−error of the estimation of the traffic flow volumes over a period of 20 time steps, which is defined as $\|\boldsymbol{x}\|_2^{-1} \|\boldsymbol{x} - \hat{\boldsymbol{x}}\|_2$ for a vector $\boldsymbol{x}$ and its estimate $\hat{\boldsymbol{x}}$. In Figure 2, we plot the temporal evolution of this error, when we deploy Netflow on 16 routers, either selected by the greedy algorithm, or as corresponding to the largest coordinates of the design $\boldsymbol{w}$ found by (W)SCOD (20 vectors $\boldsymbol{c}$ were drawn for this experiment). We have also plotted the spatial distribution of the errors, by sorting the relative L_2−errors of the 13456 OD time series, and reporting them on a graph with the proportion of traffic on the x−axis.

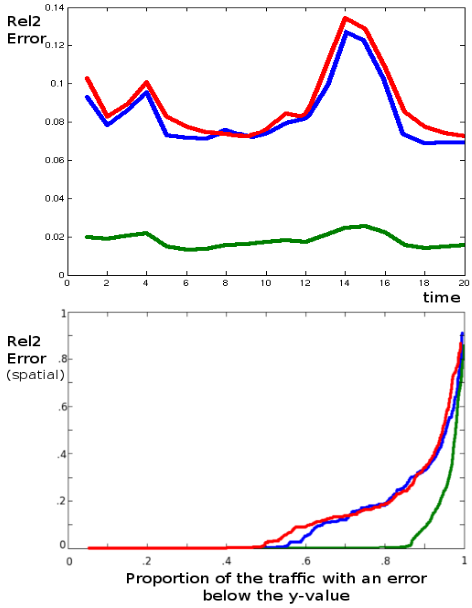

Figure 2: Temporal and Spatial L_2−Error on Opentransit, when Netflow is activated on 16 (out of 116) routers, selected by the greedy algorithm (red) or by the largest coordinates of the design found by SCOD (blue) and WSCOD (green).

In Figure 1, we plot the distribution of the L_2 spatial error for the optimal sampling problems, with the simple constraint $\sum_i w_i \leq 0.01$. Since we are not aware of any other sampling selection scheme that can be used on a network of this size, we only compare our estimate with the naive sampling rates ($w_i = 0.01/116$ for each router). For smaller instances, a comparison with the A-optimal design in a filtering context of [3] is carried out in [1].

4. CONCLUSION

In this paper, we have proposed a new method to optimize the traffic measurement, based on the estimation of a sequence of linear combinations of the flows, rather than on the estimation of the full vector of flows. This method remains tractable for very large instances, and it allows one to identify the traffic accurately.

5. REFERENCES

[1] G. Sagnol, S. Gaubert, and M.Bouhtou. Optimal monitoring in large networks by Successive **c**-Optimal Designs. Preprint available at `http://www.cmap.polytechnique.fr/~sagnol/papers/OptMonitoring_preprint.pdf`

[2] G. Sagnol. Computing optimal designs of multiresponse experiments reduces to second-order cone programming. 2009. Submitted.arXiv:0912.5467.

[3] H. Singhal and G. Michailidis. Optimal sampling in state space models with applications to network monitoring. In *Proc. ACM Sigmetrics*, pages 145–156, Annapolis, Maryland, USA, June 2–6, 2008.