Efficient Methods for Traffic Matrix Modeling and On-line Estimation in Large-Scale IP Networks

Pedro Casas*†,Sandrine Vaton*,Lionel Fillatre,‡ and Thierry Chonavel*
*TELECOM Bretagne, Brest, France - Email: name.surname@telecom-bretagne.eu
†Universidad de la República, Montevideo, Uruguay
‡Université De Technologie De Troyes, Troyes, France - Email: name.surname@utt.fr

***Abstract*—Despite a large body of literature and methods devoted to the Traffic Matrix estimation problem, the inference of traffic flows volume from aggregated data represents a key subject facing the evolution of next generation networks. This is a particular problem in large-scale carrier networks, for which efficient, accurate and stable methods for Traffic Matrix modeling and estimation are vital and challenging to conceive. In the short-term, estimation methods must be efficient and stable to allow crucial real-time tasks such as on-line traffic monitoring. In the long-term, methods must provide an accurate picture of the traffic matrix to tackle problems such as network planning, design, and dimensioning. In this paper we present and compare two efficient methods for on-line traffic matrix estimation. Based on an original parsimonious linear model for traffic flows in large-scale networks, we present a simple approach to compute an accurate traffic matrix from easily available link traffic measurements. We further extend the validation of this parsimonious model to three operational backbone networks. We analyze in depth a method to recursively estimate the traffic matrix, studying the drawbacks and omissions of the former algorithm and proposing new extensions to solve these problems. We finally perform a comparative analysis of the performance of both methods in two operational backbone networks, taking into account significant aspects such as accuracy, stability, scalability, and on-line applicability.**

***Index Terms*—Network Traffic Measurements, Modeling and Statistical Characterization, Traffic Matrix Estimation, Kalman Filtering.**

I. Introduction

Knowing and understanding the traffic that flows through a large-scale network represents a key issue in the design and engineering of the future Internet. A network-wide view of traffic flows is typically described by a traffic matrix (TM); a TM represents the volume of traffic transmitted between every pair of ingress and egress nodes of a network, also referred as the origin-destination (OD) traffic flows. The measurement of the TM is a subject of continuous debate between researchers, network operators, and technology vendors. Some of them claim that the overheads incurred in the direct measurement of the TM will become too costly and prohibitive in the future, justifying the use of aggregated data to gather its value as it has been done during the last years. This is quite a valid argument if we consider that the evolution of future access technologies and the development of optical access networks (Fiber To The Home technology) will dramatically increase the bandwidth for each end-user, triggering a brutal augmentation of the traffic to measure. Some authors forecast a value of bandwidth demand per user as high as 50 Gb/sec in 2030. On the other hand, the progress in monitoring and measurement technology of the past years make some others believe that the challenge of directly measuring the TM can be solved by improving equipment measurement capabilities and that the problem of inferring the TM from aggregated data will become obsolete. Whatever the result of this bid between increase of traffic and progress in measurement capabilities, network analysis requires efficient TM estimation methods that make use of both aggregated data and direct measurements to improve results.

Let us formally introduce the TM estimation problem. Throughout the paper, the vector $X_t = [x_t(1), .., x_t(m)]^T$ represents the value of the traffic matrix at time t, where $x_t(k)$ stands for the traffic volume of each OD flow $k = 1..m$ at time t. In a similar way, the vector $Y_t = [y_t(1), .., y_t(r)]^T$ represents the value of the links aggregated traffic volume, where $y_t(i)$ represents the total traffic volume in link $i = 1..r$ at time t. This aggregated data is available through the standard and well-known SNMP protocol, so it will be usually referred as the SNMP measurements. Given a routing matrix R, we define the traditional SNMP-TM measurement relation as:

$$Y_t = R\, X_t \tag{1}$$

where R_{ij} is equal to 1 if OD flow j traverses link i and 0 otherwise. The computation of X_t from Y_t represents a massively under constrained problem, as the number of unknown OD flows is much larger than the number of links [1]. This equation is the basis of the TM estimation problem.

A. Related Work

The problem of inferring the traffic matrix from link aggregated traffic data has been extensively studied over the past 10 years. The first approach to tackle the problem was to search for direct solutions to the ill-posed problem, introducing additional information to create additional constraints. This was achieved by TM modeling assumptions in [1], [3], deriving higher order statistics of the traffic OD flows as the additional constraints. For instance, Vardi adopted a Poisson model in [1] and Cao et al. a Gaussian model in [3]. Tebaldi et al. considered in [2] a Bayesian approach to the problem, assuming a Poisson a-priori distribution for the OD flows. A couple of years later, Medina et al. [4] showed that the basic assumptions underlying these statistical models were not justified, and that these methods performed badly when the underlying assumptions were violated.

Additional spatial information about the TM was included into the problem, taking into account the network topology and the routing process. This encouraged the application of gravity models [5] to the estimation issue. In 2003, Zhang et al [6] made a breakthrough in the TM estimation problem, combining network tomography methods [1] with gravity models to highly improve accuracy and reduce computational complexity. This method is the well-known tomogravity estimation approach. As we will show in the obtained results, the estimation performance of the tomogravity approach can be highly improved.

A final step was achieved by considering the strong diurnal patterns found in the TM [7] into the estimation problem, together with a new strong assumption not considered before: the TM can be directly measured during short periods of time. Different works were proposed in 2004 and 2005 that exploited these assumptions [8], [9], [11]. In [8] the authors proposed a pure data-driven method to estimate the TM based on the stability of the node fanouts. [7] proposed another data-driven approach to analyze OD flows, using a Principal Component Analysis (PCA) method to capture both temporal and spatial correlations. The problem with data-driven approaches is that they are highly dependent on the data they use as input and thus results can not be generalized. The last contribution was proposed in [9], [10], where a dynamic model was adopted to capture the temporal correlation of the TM, using a Kalman Filtering approach to recursively estimate the TM. These methods make use of 24hs periods of direct OD flow measurements for calibration purposes, which can be too restrictive in a future network scenario and which limits their application to many networks that currently lack measurement technology. Even more, although they seem quite accurate and they improve previous proposals, results presented in [9], [11] showed that they can be unstable and several recalibration steps should be conducted in order to provide reliable results.

B. Contributions of the Paper

In this paper we analyze two different approaches to estimate the TM from aggregated link data. The first approach considers a spatial model for OD flows previously introduced in [13] to perform an accurate TM estimation from link data. The principal virtue of this method is that it does not require direct OD flow measurements neither to perform the estimation nor to calibrate the model. We present more evidence of the relevance and applicability of this OD flow model by extending its validation to three different operational networks: the Abilene network, the GEANT network, and a Tier-2 ISP network. The second method consists of a recursive estimation of the TM, using a Kalman filtering approach as in [9]. This method makes use of direct OD flow measurements to calibrate the subjacent flow model, using then the link data to estimate the TM value. The Kalman filtering approach is quite appealing, but the original works [9], [10] present some important drawbacks and omissions we treat in this work. By introducing new simple dynamic models we enhance the performance of the approach thus improving its applicability. Both estimation algorithms are compared in terms of relevant performance indexes namely accuracy, stability, scalability, complexity, and on-line applicability among others.

The remainder of this paper is organized as follows. In section II we recall the main aspects of the linear parsimonious OD flow model, showing how it can be applied to the traffic matrix estimation problem. Section III presents and analyses different state space OD flow models for recursive estimation of the TM, analyzing the drawbacks and omissions of previous proposals. Section IV presents the evaluation of the parsimonious model for the TM estimation problem, extending its validation to three operational backbone networks. The performance of the recursive TM estimation algorithm for different OD flow models is also evaluated, both in the Abilene and GEANT backbone networks. Finally, a comparative analysis between both estimation methods and previous proposals is presented. Section IV concludes this work.

II. Parsimonious TM Modeling and TM Estimation

In this section we recall the parsimonious linear model we have previously introduced in [13], explaining how this model can be applied to tackle the TM estimation problem. The basic idea of this model is that traffic flows X_t, sorted by OD flow volume can be decomposed at each time t over a known family of q basis functions $S = \{\mathbf{s}(1), \mathbf{s}(2), \ldots, \mathbf{s}(q)\}$, with the great virtue that $q << m$ (several orders of magnitude smaller). Therefore, we assume that X_t can be expressed as:

$$X_t = S\boldsymbol{\mu}_t + \boldsymbol{\xi}_t \tag{2}$$

where $\boldsymbol{\xi}_t$ is a white Gaussian noise with covariance matrix $\Sigma = \mathrm{diag}(\sigma_1^2, \ldots, \sigma_m^2)$ that models the natural variability of the OD TM together with the modeling errors. The vector $\boldsymbol{\mu}_t = \{\mu_t(1) \ldots \mu_t(q)\}^T$ is the unknown time varying parameter vector which describes the OD flow intensity distribution with respect to the set of vectors $\mathbf{s}(i)$. We found in [13] that the order of increasing OD flows w.r.t. their traffic volume remains stable in time for several days. Figure 1 shows the OD flows traffic for (a) the Abilene network, (b) the GEANT network, and (c) a Tier-2 ISP network, sorted in the increasing order of their volume of traffic and for different time instants t. The sorted volumes of OD flows can be interpreted as a discrete non-decreasing signal with certain smoothness. The curve obtained by interpolating this discrete signal is parameterized by using a polynomial approximation. Given the shape of this curve, a cubic splines approximation is applied. A discrete spline basis is finally built, discretizing the continuous splines according to m points uniformly chosen in the interval $[1; m]$ and rearranging them according to the OD flows sorting order. The vectors $\mathbf{s}(i)$ in S correspond to the rearranged discrete splines, which form a set of basis vectors that describe the spatial distribution of the traffic. It should be clear to the reader that this model can not be generalized to all network topologies and scenarios, but that it holds for networks with a high level of traffic aggregation (e.g., a backbone network or a large

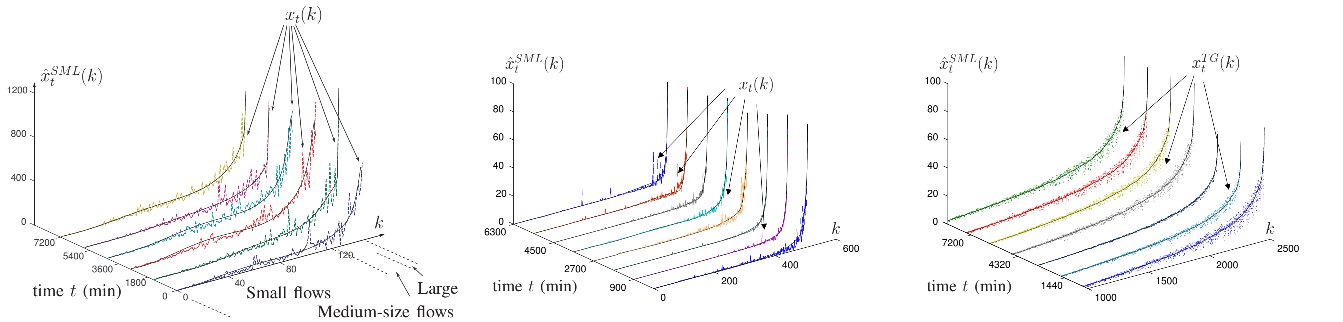

(a) Abilene Network (b) GEANT Network (c) Tier-2 ISP Network

Fig. 1. Approximation of OD flows (dashed lines) by the spline-based model (full lines) in 3 operational networks. $\hat{x}_t^{SML}(k)$ stands for the estimated OD flow k using the spline-based model, defined in equation (6). $x_t^{TG}(k)$ is the estimated OD flow k using the tomogravity estimation method, introduced in [6].

international VPN). The dashed lines in figure 1 depict the value of each sorted OD flow $x_t(k), k = 1 \ldots m$, the full lines represent the polynomial approximation of the sorted flows. In order to appreciate the time stability of this approximation, the curves are plotted for various consecutive days (at different moments of the day). Considering the SNMP-TM relation (1), the model for the link traffic is given by:

$$Y_t = G\boldsymbol{\mu}_t + \boldsymbol{\upsilon}_t, \tag{3}$$

where $G = RS$ and $\boldsymbol{\upsilon}_t \sim \mathcal{N}(0, \Phi)$, with $\Phi = R\Sigma R^T$. Since the number of columns in G is very small, the product RS and its rank can be computed very fast; therefore, we assume that G is full column rank. To simplify notation and computations, we use the whitened measurements vector Z_t:

$$Z_t = \Phi^{-\frac{1}{2}} Y_t = H\boldsymbol{\mu}_t + \gamma_t, \tag{4}$$

where $H = \Phi^{-\frac{1}{2}}\, G$, $\gamma_t \sim \mathcal{N}(0, I)$ and I is the identity matrix of correct dimensions. The purpose of this transformation is simply to whiten the Gaussian noise. Finally, the covariance matrix Σ is unknown. The solution consists of computing an estimate $\widehat{\Sigma}$ from a few measurements; in section IV-B we show that using just 1 hour of SNMP measurements is enough to provide proper results. Results on the estimation of $\widehat{\Sigma}$ can be found in [15]. This linear parsimonious model allows to estimate the OD flows volume X_t from easily available SNMP measurements Y_t. We propose to use a Maximum Likelihood (ML) estimation approach to compute an estimated traffic matrix. The ML estimate presents well established statistical properties [15]: it is asymptotically optimal, which means that it is asymptotically unbiased and efficient. Since the traffic linear model (4) is a Gaussian model, the ML estimate of $\boldsymbol{\mu}_t$, namely $\hat{\boldsymbol{\mu}}_t^{ML}$ corresponds to the least mean squares estimate:

$$\hat{\boldsymbol{\mu}}_t^{ML} = (H^T H)^{-1} H^T\, Z_t \tag{5}$$

This finally leads to the Maximum Likelihood estimate of the traffic matrix, which we will refer as the Splines-based ML (SML) estimate $\hat{X}_t^{SML}$, defined by:

$$\hat{X}_t^{SML} = S\,\hat{\boldsymbol{\mu}}_t^{ML} = \left(S(H^T H)^{-1} H^T\, \Phi^{-\frac{1}{2}} \right) Y_t \tag{6}$$

III. Recursive TM Estimation

The estimate $\hat{X}_t^{SML}$ presented in section II represents an estimation of X_t given the current value of SNMP measurements Y_t. In this section we present a method that not only uses Y_t to estimate X_t, but also takes advantage of the TM temporal correlation, using a set of past SNMP measurements $\{Y_{t-1}, Y_{t-2}, \ldots, Y_1\}$ to compute the estimate $\hat{X}_{t|t} = \mathbb{E}(X_t | Y_t, Y_{t-1}, \ldots, Y_1)$. In [9], [10], the authors use the standard Kalman filtering method [14] to recursively compute $\hat{X}_{t|t}$. We draw on the ideas of [9] as a point of departure, then analyze the weaknesses of the proposed approach, and finally extend the method to achieve more accurate and stable results.

A. A Simple State-Space Model for the Traffic Matrix

Let us consider the model that is assumed in [9], [10]. In this paper, the authors consider the traffic matrix OD flows as the hidden states of a dynamic system. A linear state space model is adopted to capture the temporal evolution of the traffic matrix, and the SNMP-TM relation (1) is used as the observation process:

$$\begin{cases} X_{t+1} &= A\, X_t + W_{t+1} \\ Y_t &= R\, X_t + V_t \end{cases} \tag{7}$$

The first equation in (7) characterizes the evolution of the OD flows X_t. A is the transition matrix that captures the dynamic behavior of the system, and W_t is an uncorrelated zero-mean Gaussian white noise that accounts both for modeling errors and randomness in the traffic flows. The second equation in (7) relates the observed links traffic Y_t to the unobserved state X_t through the routing matrix R. The measurement noise V_t is also an uncorrelated zero-mean Gaussian white noise process that models possible inconsistencies in the SNMP-TM relation. [9], [10] also assume a stationary situation where A, R, and the noise covariance matrices Q_w and Q_v are constant in time. Given this model it is possible to recursively derive the least mean squares linear estimate of X_t given $\{Y_t, Y_{t-1}, \ldots, Y_1\}$, $\hat{X}_{t|t} = \mathbb{E}(X_t | Y_t, Y_{t-1}, \ldots, Y_1)$ by using the standard Kalman filter (K-F) method. The Kalman filter is an efficient recursive filter that estimates the state X_t of a linear dynamic system from a series of noisy measurements $\{Y_t, Y_{t-1}, \ldots, Y_1\}$. It consists of two distinct phases, iteratively applied: the **Prediction Phase** uses the state estimate from the previous time-step $\hat{X}_{t|t}$ to produce an estimate of the state at the current time-step $t+1$, usually known as the "predicted" state $\hat{X}_{t+1|t} = \mathbb{E}(X_{t+1}|Y_t, Y_{t-1}, \ldots, Y_1)$,

$$\begin{cases} \hat{X}_{t+1|t} &= A\,\hat{X}_{t|t} \\ P_{t+1|t} &= A\,P_{t|t}\,A^T + Q_w \end{cases} \tag{8}$$

where $P_{t|t}$ and $P_{t+1|t}$ are the covariance matrices of the estimation error $\mathrm{e}_{t|t} = X_t - \hat{X}_{t|t}$, and the prediction error $\mathrm{e}_{t+1|t} = X_{t+1} - \hat{X}_{t+1|t}$ respectively. In the **Update Phase**, measurement at the current time-step Y_{t+1} is used to refine the prediction $\hat{X}_{t+1|t}$, computing a more accurate state estimate for the current time-step $t+1$,

$$\begin{cases} \hat{X}_{t+1|t+1} &= \hat{X}_{t+1|t} + K_{t+1}\,(Y_{t+1} - R\,\hat{X}_{t+1|t}) \\ P_{t+1|t+1} &= (I - K_{t+1}\,R)\,P_{t+1|t}\,(I - K_{t+1}\,R)^T \\ & \quad + K_{t+1}\,Q_v\,K_{t+1}^T = (I - K_{t+1}\,R)\,P_{t+1|t} \end{cases} \tag{9}$$

where K_{t+1} is the optimal Kalman gain which minimizes the mean-square error $\mathbb{E}(||\mathrm{e}_{t+1|t+1}||^2)$:

$$K_{t+1} = P_{t+1|t}\,R^T(R\,P_{t+1|t}\,R^T + Q_v)^{-1} \tag{10}$$

In order to begin the Kalman filter recursion, initial conditions $\hat{X}_{0|0}$ and $P_{0|0}$ are defined. Since the value of the initial state is unknown, the initial estimate is chosen to be $\hat{X}_{0|0} = \mathbb{E}(X_0)$ and its corresponding estimation error covariance matrix $P_{0|0} = \mathbb{E}(||\mathrm{e}_{0|0}||^2)$. The calibration of matrices A, Q_w, and Q_v requires direct OD flow measurements; in [9] the authors use a 24hs period of OD flow measurements for this purpose.

In [9], the authors adopt a non-diagonal structure to the transition matrix A, while in [10] they consider a diagonal structure to A. Both choices have major impacts when using a model like (7). If we take the expected values of the right and left hand side terms in the first equation of (7) we obtain that $m_X = A\,m_X$, where $m_X = \mathbb{E}(X_t)$ denotes the average traffic matrix value. This implies that $(I-A)m_X = 0$, that is to say that m_X should be in the kernel of $I - A$. Let us consider the case where A is a diagonal matrix. In this case, the only solution to the system $(I-A)m_X = 0$ is $m_X = 0$ and obviously this condition is not satisfied by the average traffic matrix. So particularly, the first equation in (7) is false in [10], and in this context it is only valid for centered data, i.e., $m_X = 0$. Even more, our following analysis shows that using (7) without centering the data has convergence implications. On the contrary, if we consider that A is non-diagonal, it must be calibrated in such a way that $(I-A)m_X = 0$. This is essential in the model (7) as presented in [9]. In this work the authors claim that the Kalman filter must be re-calibrated every few days, when the underlying model changes, using once again direct OD flow measurements for a new 24hs period. This seems reasonable for such a particular calibration of A. As we will show in the results, this need of recalibration can be reduced with some simple corrections to the model. Let us modify the first equation in (7) in order to have a correct state space model for the case of a diagonal state transition matrix A. If we consider the variations of the OD traffic matrix X_t around its average value m_X, i.e., $X_t^{\mathrm{c}} = X_t - m_X$, the system (7) becomes:

$$\begin{cases} X_{t+1}^{\mathrm{c}} &= A\,X_t^{\mathrm{c}} + W_{t+1} \\ Y_t &= R\,X_t^{\mathrm{c}} + V_t + R\,m_X \end{cases} \tag{11}$$

The first equation in (11) is now correct for A diagonal, which corresponds to the case of modeling the centered OD flows as spatially independent AR(1) processes; even more, the equality of expected values of the left and right hand side terms holds whatever the choice of A. In this setting the model is not as sensitive to the definition of the state transition matrix A as in (7), where the only solution is to choose A non-diagonal and such that $(I-A)m_X = 0$. However, the deterministic term that appears in the observation process violates the Kalman filter assumptions; particularly, the "measurement noise" $V_t + R\,m_X$ is not a zero-mean Gaussian process. The appropriate way of treating this problem would be to center the observation process before applying the Kalman filter, using the centered observation measurements vector $Y_t^{\mathrm{c}} = Y_t - \mathbb{E}(Y_t) = Y_t - R\,m_X$. Nevertheless, we apply the Kalman filter equations to system (11) in order to appreciate the impact of using non-centered observation data when A is diagonal. Let us define $\tilde{X}_{t|t}$ as the estimate that one would obtain if the Kalman equations (8) and (9) were applied with the non-centered SNMP measurements Y_t as input. Using the Kalman filter equations, we can express both the evolution of the estimate $\hat{X}_{t|t}^{\mathrm{c}} = \mathbb{E}(X_t^{\mathrm{c}}|Y_t^{\mathrm{c}}, \ldots, Y_1^{\mathrm{c}})$ and the evolution of $\tilde{X}_{t|t}$ as:

$$\begin{matrix} (*) & \hat{X}_{t+1|t+1}^{\mathrm{c}} &= A\,\hat{X}_{t|t}^{\mathrm{c}} + K_{t+1}\,(Y_{t+1}^{\mathrm{c}} - RA\,\hat{X}_{t|t}^{\mathrm{c}}) \\ (**) & \tilde{X}_{t+1|t+1} &= A\,\tilde{X}_{t|t} + K_{t+1}\,(Y_{t+1} - RA\,\tilde{X}_{t|t}) \end{matrix} \tag{12}$$

where we have assumed the same Kalman gain in both equations as its value does not depend on the observations. If we define the error $\eta_t = \tilde{X}_{t|t} - \hat{X}_{t|t}^{\mathrm{c}}$, the difference between $(**)$ and $(*)$ can be written as:

$$\eta_{t+1} = (I - K_{t+1}\,R)A\,\eta_t + K_{t+1}\,R\,m_X \tag{13}$$

Let us assume that the Kalman filter converges; in that case, we can substitute the Kalman gain in (14) by its limit value $K = \lim_{t\to\infty} K_t$:

$$\eta_{t+1} = (I - KR)A\,\eta_t + KR\,m_X \tag{14}$$

Without loss of generality, let us suppose that $\eta_0 = 0$. We are going to prove that an error term is propagated and that the error either diverges to infinity or converges to a constant non-null value. As $\eta_0 = 0$, we can express η_t as:

$$\eta_t = \sum_{k=0}^{t-1} ((I - KR)A)^k KR\,m_X, \quad \forall t > 0 \tag{15}$$

If the spectral radius of $(I - KR)A$ is greater than 1, then the error term η_t diverges to infinity. On the contrary, if the spectral radius of $(I - KR)A$ is lower than 1, then the error term η_t converges to a constant value:

$$\eta_\infty = \lim_{t\to\infty} \eta_t = (I - (I - KR)A)^{-1} KR\,m_X \tag{16}$$

This shows that, when considering a diagonal structure for the state transition matrix A in (7), not only the state space model is false but even after centering the data and explicitly introducing the mean value m_X, the Kalman filter does not converge to the real value of the traffic matrix if non-centered data Y_t is used in the filter. On the contrary, there is a gap between the real an the estimated value that is proportional to m_X (this is verified in the results in section IV-C).

B. State-Space model for centered TM variations: static mean

This problem can be easily solved in different ways. As we said, the most obvious solution would be to consider a centered observation process Y_t^c. However, we will consider a more standard approach: a deterministic term in the observation process can always be removed by adding a new deterministic state to the state model. Let us define a new state variable $U_t = [X_t^c \; m_X]^T$. In this case, (11) becomes:

$$\begin{cases} U_{t+1} &= \begin{bmatrix} A & \mathrm{O} \\ \mathrm{O} & I \end{bmatrix} U_t + \begin{bmatrix} W_{t+1} \\ \mathrm{O} \end{bmatrix} = C\,U_t + \Psi_{t+1} \\ Y_t &= \begin{bmatrix} R & R \end{bmatrix} U_t + V_t \;=\; B\,U_t + V_t \end{cases} \tag{17}$$

where O is the null matrix of accurate size. This new model has twice the number of states, augmenting the computation time and complexity of the Kalman filter. However, it presents several advantages: (i) it is not necessary to center the observations Y_t; (ii) the matrix A can be chosen as a diagonal matrix, which corresponds to the case of modeling the centered OD flows as AR(1) processes. Autoregressive models have been widely applied in the traffic matrix literature [12]; as we show in the results, obtained results with a simple AR(1) model and the K-F technique are accurate compared to the target error for standard traffic matrix estimation tools (about 10% [9], [11]) and this is clearly much easier and more stable than calibrating a non-diagonal matrix such that $(I - A)m_X = 0$; in fact, authors in [10] observe that re-calibrations are often not needed when using a diagonal transition matrix, and the results we obtain are stable during the whole evaluation period of 1 week, which is not the case in [9]; (iii) the Kalman filter estimates the mean value of the OD flows m_X, assumed constant in (17), and finally (iv) this model allows to impose a dynamic behavior to m_X, improving the estimation properties of the filter. This is exactly the step we take next.

C. Extending the model: dynamic mean

Using model (17) with the Kalman filtering technique produces quite good estimation results as we show in section IV-C. However, this model presents a major drawback: it assumes that the mean value of the OD flows m_X is constant in time. We improve (17) by adopting a simple dynamic model for m_X, in order to allow small variations of the OD flows mean value:

$$m_X(t+1) \;=\; m_X(t) + \boldsymbol{\zeta}_{t+1} \tag{18}$$

where $m_X(t)$ represents the dynamic mean value of X_t and $\boldsymbol{\zeta}_t$ is a zero-mean white Gaussian noise process with covariance matrix $Q_{\boldsymbol{\zeta}}$. This model corresponds to a random walk process, which is commonly applied to describe several dynamic models in economics, physics, etc. In this context, (17) becomes:

$$\begin{cases} U_{t+1} &= \begin{bmatrix} A & \mathrm{O} \\ \mathrm{O} & I \end{bmatrix} U_t + \begin{bmatrix} W_{t+1} \\ \boldsymbol{\zeta}_{t+1} \end{bmatrix} = C\,U_t + \Theta_{t+1} \\ Y_t &= \begin{bmatrix} R & R \end{bmatrix} U_t + V_t \;=\; B\,U_t + V_t \end{cases} \tag{19}$$

As we see in the results in section IV-C, such a simple model provides more accurate and more stable results.

IV. Evaluation and Discussion

In this section we present the evaluation of the estimation algorithms using real measurements from different operational backbone networks. We first describe the datasets used in the evaluation, then evaluate the SML estimation method and extend the validation of the splines-based model, then we evaluate the recursive Kalman filter estimation technique for the different proposed state-space models, and finally we present a comparative analysis of both algorithms.

A. The Datasets

Network	n^o nodes - links	n^o ODFlows	Data	Sampling
Abilene	12 - 54	132	OD flows traffic	5’
GEANT	23 - 74	506	OD flows traffic	15’
Tier-2 ISP	50 - 168	2450	links traffic	10’

TABLE I
Network Topologies for the Datasets.

The evaluation of the estimation algorithms is conducted using real data from two operational networks: the Abilene network, an Internet2 backbone network, and the GEANT network, a European research network. For the validation of the splines-model, we also include data from a private Tier-2 ISP network. Table I presents the topology of each network. Abilene traffic data consists of 5’ sampled TMs collected via Netflow from the Abilene Observatory [17] and available at [18]. GEANT traffic data consists of 15’ sampled TMs, built from IGP and BGP routing information and Netflow data in [19], available on the TOTEM website [20]. The Tier-2 ISP network is a private network and data is not public. Direct OD flow measurements are not available for this network. Instead, link traffic volumes are gathered each 10’ via SNMP. Using this data and a rich description of the topology, we perform a tomogravity estimation [6] of the real OD flows volume. The tomogravity method is a widely accepted method to estimate OD flow volumes from link traffic measurements and topology information with confident results. In the numerical validation of our splines model for the Tier-2 ISP network, we show that the obtained estimation results are very close to those obtained with the tomogravity estimate for this network. In the following evaluations, we assume that traffic flows X_t are just known during the calibration of the recursive Kalman algorithm and consider the SNMP measurements Y_t as the input known data. In order to verify the stability properties of the proposed models, two sets of measurements are used for each network topology: the "learning" dataset, used for calibration purposes, and the "testing" dataset, used to evaluate the performance of the algorithms. Let T_{learn} and T_{test} be the sets of time indexes associated with measurements from the learning and testing datasets respectively.

B. Validation of the Splines Model and SML TM Estimation

In this case, both the learning and testing datasets consists of SNMP measurements. The learning dataset is composed of one hour of SNMP measurements and it is used to construct the splines basis S; the testing dataset is composed of 672 SNMP measurements. The splines-based model is computed for each network using each learning dataset, following these steps: (i) the tomogravity (TG) estimate $\hat{x}_t^{\,TG}(k)$ is computed for all OD flows k and all $t \in T_{\text{learn}}$; (ii) the mean flow values $\bar{x}^{\;TG}(k) = \frac{1}{\#(T_{\text{learn}})} \sum_{t \in T_{\text{learn}}} \hat{x}_t^{\,TG}(k)$ are computed, where $\#(T_{\text{learn}})$ is the number of time indexes in the learning dataset; (iii) finally, the obtained mean values $\bar{x}^{\;TG}(k)$ are sorted in ascending order to obtain a rough estimate of the OD flows traffic volume. The spline-based model is designed with cubic splines and 2 knots, representing small, medium-size, and large OD flows. The mean value $\bar{x}^{\;TG}(k)$ of each OD flow is used to compute an estimate $\hat{\sigma}_k^2$ of σ_k^2, which leads to an estimate $\hat{\Phi}$ of Φ, quite efficient and sufficient in practice.

As a global indication of the accuracy of the SML estimate and to test the performance of the short learning step, we apply the relative root mean squared error (RRMSE) for each time t in the testing dataset:

$$\text{RRMSE}(t) = \frac{\sqrt{\sum_{k=1}^{m} \left(x_t(k) - \hat{x}_t^{SML}(k)\right)^2}}{\sqrt{\sum_{k=1}^{m} x_t(k)^2}}, \;\; \forall t \in T_{\text{test}} \quad (20)$$

where $x_t(k)$ is the true traffic volume of OD flow k at time t and $\hat{x}_t^{SML}(k)$ denotes the corresponding SML estimate previously defined in (6). The RRMSE provides at each time t a summary of the relative estimation error for all m OD flows. In the validation of the model for the Tier-2 ISP network, we compare the value of the SML estimate $\hat{x}_t^{SML}(k)$ against the tomogravity estimate $\hat{x}_t^{TG}(k)$, using the relative root mean squared difference (RRMSD) between both estimates:

$$\text{RRMSD}(t) = \frac{\sqrt{\sum_{k \in \textit{topTG-}T_h} \left(\hat{x}_t^{TG}(k) - \hat{x}_t^{SML}(k)\right)^2}}{\sqrt{\sum_{k \in \textit{topTG-}T_h} \left(\hat{x}_t^{TG}(k)\right)^2}}, \;\; \forall t \in T_{\text{test}} \quad (21)$$

Comparing all flows in (20) is not a reasonable approach. The tomogravity estimate provides quite accurate results for relatively high-volume flows, but poor for small flows [6]; we define the *topTG-*T_h flows as those estimated flows by the tomogravity method that are reasonably stable in time and which mean value exceeds a threshold T_h. In this sense we only keep the most accurately estimated flows, removing the noisy or erratic estimates which seems to be wrongly estimated. Figure 2.(a) presents the temporal evolution of the RRMSE for the 672 measurements in the testing datasets for Abilene and GEANT. In both cases, the relative error remains stable in time, reinforcing the observations about time-stability of the model we drew from figure 1. Figure 2.(b) shows that more than 70% of the time, estimation relative errors are below 10%. A deeper study of the RRMSE shows that in most cases, large RRMSE values correspond to large relative errors in the lowest-volume OD flows, which are well known to be hard to estimate [6], [11]. Note however that small OD flows have little impact on traffic engineering tasks and so are generally less important to estimate. The mean values of the RRMSE for the evaluation period are 8.14% for Abilene and 7.04% for GEANT. Methods proposed in the literature as "accurate" estimates present relative errors that vary between 5% and 15% [9], [11], so obtained results are satisfactory. Figure 3 depicts the temporal evolution of the RRMSD between the TG and SML estimates, for a Tier-2 ISP network. In this evaluation, we tune T_h such that 60% of the total flows are compared in the RRMSD index, which represents approximately 95% of the total traffic. The relative difference between the TG and the SML estimates is stable in time and has a mean value of 0.57%. Based on our previous observations about the tomogravity estimate, we conclude that the splines model is also accurate for this Tier-2 ISP network. As a final validation of the splines-model, we verify the Gaussian assumption for Abilene and GEANT. The "residuals" of measurements are analyzed, i.e., the obtained traffic after filtering the mean part $H\boldsymbol{\mu}_t$. The residuals are obtained by projection of the whitened measurements vector Z_t onto the left null space of H. Quantile-Quantile plots for two of these residual processes are plotted in figure 4, both for Abilene and GEANT. These residual processes clearly

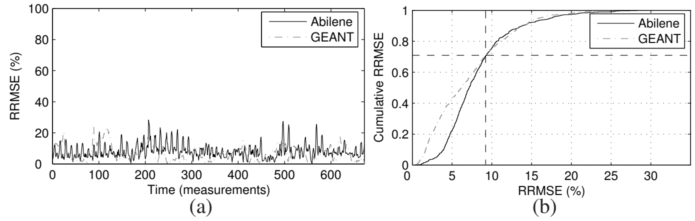

Fig. 2. (a) RRMSE(t) and (b) Cumulative RRMSE(t) for 672 measurements in Abilene and GEANT

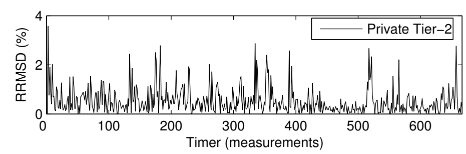

Fig. 3. RRMSD(t) for 1500 flows in a Tier-2 ISP network

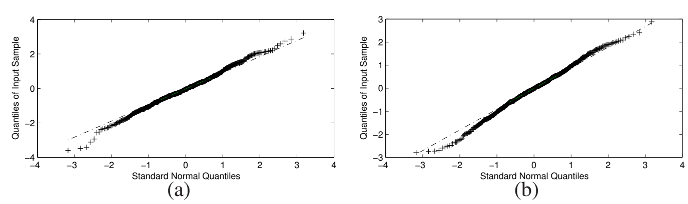

Fig. 4. QQ-plots for 2 residual processes from (a) Abilene and (b) GEANT.

follow a Gaussian distribution. We also verify the Gaussian assumption by applying a Kolmogorov-Smirnov goodness-of-fit hypothesis test to the residual processes. The acceptance rate of this test at the level 5% is 98.5% for Abilene and 97.7% for GEANT, which also confirms the Gaussian assumption.

C. Recursive TM Estimation

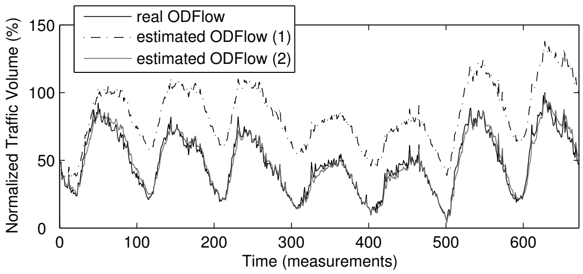

Fig. 5. Estimated OD flows using K-F for (1) model 7 and (2) model 17

The first evaluation consists of evidencing the convergence problem of the recursive TM estimation when using a model like (7) with A diagonal, as it is done in [10]. In this sense, we compare the performance of the Kalman filter using models (7) and (17). In both cases we adopt a diagonal structure for the state transition matrix A, namely an AR(1) model for each OD flow. In this evaluation and through the rest of section IV-C, the learning dataset is composed of 24hs of direct OD flow measurements X_t, as it is the case in [9]. The testing dataset consists of 1 week of SNMP measurements from the GEANT dataset, which represents 672 measurements. We also assume that the relation between X_t and Y_t is exact, that is to say $V_t = 0, \forall t$. The learning dataset is used to calibrate both models (7) and (17), namely estimate the corresponding transitions matrices and noise covariance matrices (the AR(1) parameters). We use the Yule-Walker method to compute these matrices. This method solves the Yule-Walker equations for the AR processes by means of the Levinson-Durbin recursion, see [16] for details. Figure 5 depicts the estimation of one sample OD flow with both Kalman filters; the full black curve represents the real OD flow; the dashed black curve depicts the estimated OD flow using model (7); the full gray curve depicts the estimated OD flow using model (17). In both cases, the Kalman filter properly tracks the real traffic pattern, as both curves shape are similar to the real one. However, there is a clear error-gap when using model (7), which comes from our previous analysis. Figure 6 shows the evolution of the relative

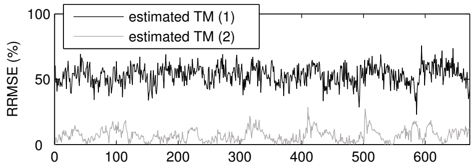

Fig. 6. RRMSE(t) for (1) model 7 and (2) model 17

estimation error RRMSD(t) for all OD flows of the TM. The mean relative error is 53.4% for model (7) and 6.2% for model (17); in both cases the error evolution is quite stable around its mean value during the whole evaluation week, giving a first evidence of the stability advantages of a diagonal transition matrix.

We now compare the estimation performance of the Kalman filter for models (17) and (19), namely assuming a constant mean value for OD flows or a random walk process, and a diagonal transition matrix in both cases. For this purpose, we consider a week of traffic in Abilene and GEANT. We consider the same assumptions adopted in the previous evaluation and calibrate the different matrices in the same way. In order to estimate the covariance matrix $Q_{\boldsymbol{\zeta}}$ of the random walk noise process $\boldsymbol{\zeta}_t$, we take the following steps: using a sliding window averaging filter we first remove the fast temporal variations from the direct OD flow measurements of the learning dataset. For each OD flow time-series, we consider the approximate derivative time-series (i.e., the difference of consecutive measurements) and compute its variance. We finally use this variance as an estimate of each diagonal element in $Q_{\boldsymbol{\zeta}}$. Figure 7 depicts the relative estimation error

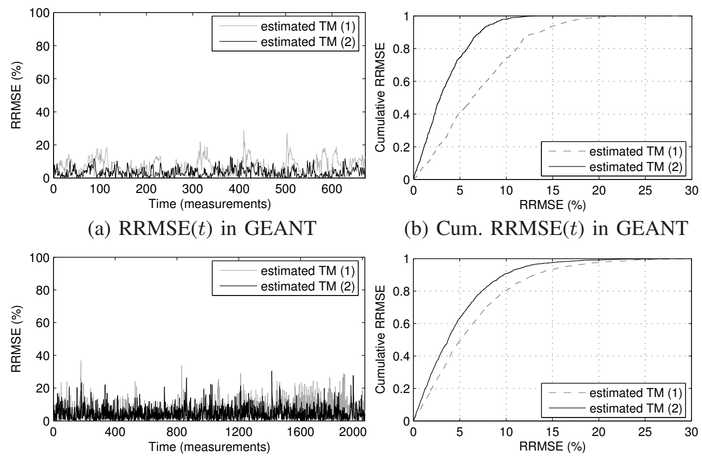

(a) RRMSE(t) in GEANT (b) Cum. RRMSE(t) in GEANT

(c) RRMSE(t) in Abilene (d) Cum. RRMSE(t) in Abilene

Fig. 7. RRMSE(t) and Cumulative RRMSE(t) for 1 week of traffic in GEANT and Abilene, using (1) model 17 and (2) model 19

evolution for all TM OD flows using both models and one week of measurements in GEANT and Abilene. The cumulative RRMSE is also depicted in these figures. The obtained mean values of the relative errors are 6.20% and 4.23% in GEANT and 6.87% and 4.48% in Abilene, for models (17) and (19) respectively. We can draw two important conclusions from both evaluations: in both cases, considering a variable mean value $m_X(t)$ produces better results, both as regards accuracy and stability, as the curve of cumulative RRMSE shows a sharper growth. The second conclusion is about the advantage of correctly using a diagonal transition matrix; in all evaluations the stable evolution of the error shows that the underlying model remains valid during several days when considering such a transition matrix, a major advantage with respect to the results obtained in the former work [9]. This simple observation has a major impact on the applicability of the method in a real scenario: if the underlying model remains stable, it is not necessary to conduct periodical re-calibrations, dramatically reducing measurement overheads.

D. Comparative Analysis and Discussion

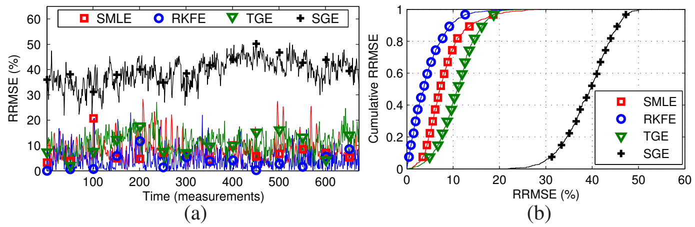

Fig. 8. (a) RRMSE(t) and (b) Cumulative RRMSE(t) for 672 measurements in Abilene, for the SMLE, the RKFE, the TGE, and the SGE.

Figure 8 presents a comparative summary of the performance of both presented methods in Abilene. The Splines-Based Maximum Likelihood Estimate (SMLE) and the Recursive Kalman Filter Estimate (RKFE) are compared against two very well known traffic matrix estimation algorithms used as baseline: the Simple Gravity Estimation method (SGE) and the Tomo-Gravity Estimation method (TGE). The obtained mean values of the relative error are 8.14%, 4.48%, 11.15%, and 39.08% for the SMLE, RKFE, TGE, and SGE respectively. From figure 8.(b) we can see that the SMLE and the RKFE produce estimation relative errors below 10% for approximately 75% and 92% of the TMs respectively, while this result drops to nearly 40% for the TGE, and to 0% for the SGE. These results allow to show the improvements of both proposed algorithms w.r.t. previous highly respected work.

To conclude with the evaluation, we present in table II a comparative analysis between the SMLE and the RKFE methods. Let us discuss each of the compared items. As regards accuracy, the RKFE presents better results, which is quite evident given its use of past data to compute the current estimate. The learning data used by the SMLE consists of pure SNMP measurements, and the method uses a remarkably short learning step. In this sense, the SMLE can be applied in networks where direct OD flow monitoring technology is not available. On the contrary, a 24hs period of direct OD flow measurements is needed to calibrate the RKFE method. As regards complexity, both algorithms are simple to implement and calibrate, specially the RKFE after the modifications introduced in this work. The considered assumptions in deriving the SMLE are quite strong compared to those adopted by the RKFE method. Nevertheless, the validation of the splines model in three different networks shows that these assumptions are correctly verified in these cases. Thanks to the underlying parsimonious model adopted in the SMLE, the method is completely scalable with the size of the network. The RKFE method does not scale with the number of OD flows in terms of computational time and memory issues, given its intrinsic recurrent characteristic and the inversion of large matrices. The scalability problem can be alleviated by implementing faster pseudo-inversion algorithms, but the problem still remains. Both algorithms can be directly applied for on-line tasks such as traffic monitoring, but the RKFE presents an interesting advantage, namely the ability to predict future values of the TM, taking advantage of the strong temporal correlation of OD flows traffic. The short learning step of the SMLE method allows its use under dynamic routing conditions, provided that the routing modifications occur at time intervals longer than 1 hour in order to allow a correct model recalibration.

Performance Index	SMLE $\hat{X}_t^{SML}$	RKFE $\hat{X}_{t\|t}$
Mean RRMSE (%) - Abilene	8.14	4.48
Mean RRMSE (%) - GEANT	7.04	4.23
Learning Data/Input Data	SNMP/SNMP	TM/SNMP
Learning Period Duration	1 hs.	24 hs.
Complexity	simple	simple
Assumptions	strong	relatively weak
Scalability	yes	partial
On-line Computation	yes	yes
Prediction Enable	no	yes
Supports Dynamic Routing	partially	no

TABLE II
COMPARATIVE PERFORMANCE OF THE SMLE AND THE RKFE.

V. CONCLUSIONS

In this paper we have revisited the TM estimation problem, dealing with important issues such as accuracy, stability, scalability, and on-line applicability among others. We have extended the validation of a previously introduced spatial model for OD flows in a large-scale network to three different operational backbone networks and showed how this model can be efficiently used in the TM estimation problem. We have introduced a simple state space model for OD flows and use it to recursively estimate the TM, using a Kalman filtering approach. A deep analysis of the drawbacks of this method as it was originally introduced allowed to better understand the stability problems of the original approach and to propose simple yet effective improvements. Both algorithms present better results than the most accepted estimation methods in the field.

REFERENCES

[1] Y. Vardi, "Network tomography: estimating source-destination traffic intensities from link data", in *J. Amer. Statist. Assoc*, 91, pp. 365-377, 1996.
[2] C. Tebaldi et al, "Bayesian inference on network traffic using link count data", in *J. Amer. Statist. Assoc*, 93, pp. 557-576, 1998.
[3] J. Cao et al, "Time-varying network tomography", in *J. Amer. Statist. Assoc*, 95, pp. 1063-1075, 2000.
[4] A. Medina, K. Salamatian, S. Bhattacharyya and C. Diot, "Traffic Matrix Estimation: Existing Techniques and New Directions", in *Proc. ACM SIGCOMM*, 2002.
[5] J. Kowalski and B. Warfield, "Modeling traffic demand between nodes in a telecommunications network", in *ATNAC*, 1995.
[6] Y. Zhang et al, "Fast Accurate Computation of Large-Scale IP Traffic Matrices from Link Load Measurements", in *Proc. ACM SIGMETRICS*, 2003.
[7] A. Lakhina et al, "Structural Analysis of Network Traffic Flows", in *Proc. ACM SIGMETRICS*, 2004.
[8] K. Papagiannaki, et al, "A Distributed Approach to Measure Traffic Matrices", in *Proc. ACM IMC*, 2004.
[9] A. Soule et al."Traffic Matrix Tracking using Kalman Filters", in *LSNI*, 2005.
[10] A. Soule, K. Salamatian and N. Taft, "Combining Filtering and Statistical Methods for Anomaly Detection", in *Proc. USENIX/ACM IMC*, 2005.
[11] A. Soule et al, "Traffic Matrices: Balancing Measurements, Inference and Modeling", in *Proc. ACM SIGMETRICS*, 2005.
[12] Y. Zhang et al, "Network Anomography", in *Proc. USENIX/ACM IMC*, 2005.
[13] P. Casas, L. Fillatre and S. Vaton, "Robust and Reactive Traffic Engineering for Dynamic Traffic Demands", in *Proc. EuroNGI*, 2008.
[14] M. Hayes, "Statistical Digital Signal Processing and Modeling", J. Wiley & Sons, 1996.
[15] C. Rao, "Linear Statistical Inference and its Applications", J. Wiley & Sons, 1973.
[16] S. Kay, "Modern Spectral Estimation: Theory and Appl", Prentice-Hall, 1988.
[17] The Abilene Observatory, *http://abilene.internet2.edu/observatory/*
[18] Y. Zhang, "Abilene Dataset 04", *http://www.cs.utexas.edu/∼yzhang/*
[19] S. Uhlig et al, "Providing Public Intradomain Traffic Matrices to the Research Community", in ACM Sigcomm Computer Communication Review, 2006.
[20] TOTEM Traffic Engineering Toolbox, *http://totem.run.montefiore.ulg.ac.be/*