Optimal Sampling in State Space Models with Applications to Network Monitoring

Harsh Singhal
Department of Statistics
The University of Michigan
singhal@umich.edu

George Michailidis
Department of Statistics and
Electrical Engineering & Computer Science
The University of Michigan
gmichail@umich.edu

ABSTRACT

Advances in networking technology have enabled network engineers to use sampled data from routers to estimate network flow volumes and track them over time. However, low sampling rates result in large noise in traffic volume estimates. We propose to combine data on individual flows obtained from sampling with highly aggregate data obtained from SNMP measurements (similar to those used in network tomography) for the tracking problem at hand. Specifically, we introduce a linearized state space model for the estimation of network traffic flow volumes from combined SNMP and sampled data. Further, we formulate the problem of obtaining optimal sampling rates under router resource constraints as an experiment design problem. Theoretically it corresponds to the problem of optimal design for estimation of conditional means for state space models and we present the associated convex programs for a simple approach to it. The usefulness of the approach in the context of network monitoring is illustrated through an extensive numerical study.

Categories and Subject Descriptors

C.2.3 [**Computer-Communication Networks**]: Network Operations— *Network Monitoring*

General Terms

Algorithms, Design, Measurement, Theory

Keywords

Internet Traffic Matrix Estimation, Kalman Filtering, State Space Models, Optimal Design of Experiments

1. INTRODUCTION

Monitoring plays an important role in network management tasks, such as capacity planning by tracking demands and forecasting traffic, identifying failures together with their causes and impact, detecting malicious activity and configuring routing protocols [1, 21]. Traditionally, network monitoring has focused on single links, but that is proving insufficient in today's complex networks [12]. A network-wide view of traffic is essential for many management tasks and hence monitoring systems with such capabilities are particularly useful. However, monitoring everywhere is expensive and thus judicious collection of limited measurements combined with modeling techniques proves crucial [17].

There has been considerable work on the problem of estimation and monitoring of flow volumes from aggregate SNMP data [18]. Indeed, the whole field of network tomography addresses this problem. Specifically, models to accomplish this goal have been discussed in [3, 13, 25, 10], identifiability issues examined in [22, 4, 20], while tracking flow volumes studied in [21]. In the latter paper, a state space model is proposed for flow volumes which are in turn tracked using SNMP and Kalman filtering. However, one issue is that this approach heavily relies on low dimensional projections of flow volumes and no information from the orthogonal subspace is available.

On the other hand, direct measurements of the flows traversing the network simplify the monitoring task significantly. However, resource limitations (cache memory, required bandwidth, etc) and the high volume of traffic in today's high speed network allow only for limited measurements through *sampling*. Traffic is carried on packets that can be sampled at router interfaces, henceforth called *observation points*. Each packet from the aggregate flow at an observation point is sampled independently with a certain probability (sampling rate). Typical sampling rates are about .01. For every packet sampled, its header information is recorded which allows one to reconstruct objects of interest, such as volumes of flows with a particular source and destination traversing the network. An important issue is how to *select* (design) the sampling rates across the network subject to resource constraints to in order to collect the maximum amount of information on the underlying source-destination flows.

The use of sampled data in networking has become an active area of research. There has been a fair amount of work on ways to augment SNMP data with small amount of sampled data when necessary [15]. The focus of some of the current research is on simultaneously controlling volume and accuracy of such data [7]. For example, Choi et. al. [5] discuss some of the considerations regarding sampling error and measurement overhead for some simple sampling

SIGMETRICS'08, June 2–6, 2008, Annapolis, Maryland, USA.

schemes. Duffield et. al. [9] investigate this issue for different sampling schemes including threshold sampling, uniform flow sampling, uniform packet sampling and sample and hold. In addition they provide algorithms for dynamic control of sample volume. Another interesting area is the optimal combination of sampled data from across the network [9, 8]. Yang et al. [23, 24] study estimation of individual flow distributions through non-parametric techniques based on sampled data. Zhao et. al. [27] investigated the problem of combining (possibly dirty) SNMP and sampled data.

The objective of this paper is to design optimal sampling rates simultaneously for all observation points, which combined with aggregate SNMP measurements would allow one to track flow volumes over time. The main contributions of the paper are the following. First, we motivate the need for combining aggregate SNMP data, together with directed sampled flow measurements from across the network and time intervals in order to obtain better accuracy in flow monitoring. This is accomplished through a linear state space model. Based on this model we formulate the problem of optimally designing (selecting) the sampling rates under resource constraints. Second, we propose the general problem of optimal design for estimation of conditional means for state space models and present the associated convex programs for a greedy approach. This provides a mathematical framework for the task at hand. Finally, we undertake an extensive numerical investigation of optimal sampling schemes for estimation of flow volumes for reasonable values of associated parameters and under model misspecification.

The remainder of the paper is organized as follows. In section 2, we motivate the need for combining data over time and from across the network. Further, we present spatio-temporal models for flow volumes and measurement data. In section 3, we formulate the optimal sampling problem as optimal experiment design problem and derive specific representations of the covariance of estimation error. In section 4, we use the general form of the optimization problem to present the convex programs associated with these problems. In the first part of section 5 we present performance evaluation of the proposed methodology when modeling assumptions are met. In the second part, we show that some level of misspecification is inevitable and demonstrate the effects of these on the performance of the proposed methodology. We conclude in section 6 with remarks on the present and future work.

2. FLOW AND MEASUREMENT MODELING

Consider a network comprised of nodes and links. A flow traversing the network is determined by its source and destination nodes and its path across the network is considered known and pre-specified (deterministic routing). Suppose that there are n_r such flows and one can obtain accurate measurements (e.g. SNMP) about the sum of the flows at the n_l links. Typically, the number of flows is significantly larger than the number of links; i.e. $n_l << n_r$.

Let $X(t) = (X_1(t), \cdots, X_{n_r}(t))'$ be the vector of a random realization of the quantities of interest, such as volumes of different flows, in time interval t. Further, assume that n_l SNMP measurements $Y(t) = (Y_1(t), \cdots, Y_{n_l}(t))'$ for the same time interval t are available. Each SNMP measurement is an estimate of the sum of all traffic traversing a particular link in a particular measurement interval. Thus, we can write

$$Y(t) = AX(t) + u(t)$$

where A is a $n_l \times n_r$ routing 0/1 matrix that describes the routes of the various network flows and $u(t)$ is independent noise. It is assumed that $Cov(u(t)) = \sigma^2 I$. We assume that the matrix A does not change over time. However, the case of time varying routing is handled in a straightforward way.

Notice that direct estimation of $X(t)$ is not feasible, since the system is ill-conditioned ($n_l < n_r$).

Availability of additional sampled data on *individual* flows should clearly improve our ability to track $X(t)$. However, such data can be fairly noisy at low sampling rates. As an illustrative example, suppose that a flow with volume X in a certain time interval is sampled at a rate ξ. If the number of sampled packets is N, then the usual [8] estimate of flow volume is $Z \equiv N/\xi$. Now if X is distributed as a Poisson random variable and the conditional distribution of N given X is assumed to be binomial with parameters X and ξ, then it can easily be shown that correlation between X and Z is equal to $\sqrt{\xi}$. This can be quite low for realistic sampling rates in networks. Figure 1 shows a scatter plot of 100 independent and identically distributed (i.i.d) samples of (X, Z) pairs. This example strongly suggests that flow volumes could be tracked more accurately by combining sampled data from across the network and (more crucially) across measurement intervals.

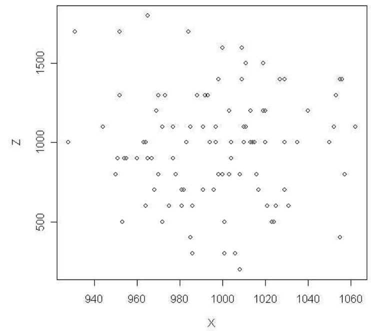

Figure 1: Sampling noise in estimated value Z for sampling rate $\xi = .01$

Next, assume that in addition to aggregate measurements $Y(t)$, we have the ability to obtain direct measurements $Z(t)$. The key feature of these measurements is that their covariance is influenced by the measurement design. We will assume a linear model:

$$Z(t) = LX(t) + \epsilon(t)$$

where L is a known $n_g \times n_r$ matrix of full column rank and $\epsilon(t)$ is independent of $X(t)$ with

$$\text{Cov}(\epsilon(t)) = \Psi(t, \xi(t))^{-1}$$

where $\Psi(t, \cdot)$ is a linear function and $\xi(t)$ is the value of design variables in time interval t. We will focus on the case when $\Psi(t, \xi(t))$ is a diagonal matrix. Generally this matrix would be singular if one or more elements of $\xi(t)$ are 0. For a given $\xi(t)$ the problem can be redefined such that the new covariance matrix is non-singular. However, the optimization problem is not affected since it depends only on $L'\Psi(t,\xi(t))L$ as will be shown later on.

Notice that this linear model approximates the setup for estimating flow volumes from sampled traffic. Suppose there are n_o observation points on the network where direct measurements on the flows can be collected. These may be routers or router interfaces. Further, assume that sampling rates of $\xi(t) = (\xi_1(t), \cdots, \xi_{n_o}(t))'$ are used at observation points $1, \cdots, n_o$, respectively, during interval t. Any given observation point $k \in \{1, \cdots, n_o\}$ generates estimates for g_k elements of X_t; i.e. the number of flows that go through that node. Thus, a total of $n_g = \sum_{k=1}^{n_o} g_k$ measurements are available, which need to be optimally combined to get the required estimates. In this paper, we focus exclusively on the uniform packet sampling scheme.

Next, consider a single observation point k with sampling rate ξ_k in a particular time interval and assume that the set $\mathcal{C}_k$ of flows traverse it. Each packet sampled from a flow is recorded by its corresponding counter. Assume that the i-th counter records N_i, the number of packets sampled from flow $l(i)$ in that time interval. We can write

$$N_i = \sum_{j=1}^{X_{l(i)}} W_{ij},$$

where $X_{l(i)}$ is the volume of traffic in flow $l(i)$ in the same time interval and W_{ij} are i.i.d. Bernoulli random variables with mean ξ_k for $i \in \mathcal{C}_k$. Now,

$$E[N_i|X] = X_{l(i)}E[W_{ij}] = X_{l(i)}\xi_k.$$

Thus $Z_i \equiv N_i/\xi_k$ is an unbiased estimator of $X_{l(i)}$. Following [6] we can write

$$\begin{aligned} Cov(Z_i|X) &= X_{l(i)}Cov(W_{ij})/\xi_k^2 \\ &= X_{l(i)}\xi_k(1-\xi_k)/\xi_k^2 \simeq X_{l(i)}/\xi_k \end{aligned} \quad (1)$$

Thus, in vector notation we get

$$E[Z|X] = LX$$

where $L_{ij} = 1$ only if $l(i) = j$ and 0 otherwise and $Cov(Z|X) \simeq D$ where D is a diagonal matrix. Using (1), the inverse of D is given by $D^{-1} = \sum_k \xi_k \Psi_k$, where Ψ_k are diagonal matrices with

$$[\Psi_k]_{ii} = \begin{cases} 1/X_{l(i)} & \text{if } i \in \mathcal{C}_k \\ 0 & \text{otherwise} \end{cases} \quad (2)$$

Clearly the matrices Ψ_k depend on the unknown X. However, we will assume that we can use some predetermined matrix as approximation in its place. This is a standard assumption in locally optimal design methodology for nonlinear models [19]. We will examine the effects of such a misspecification numerically in Section 5.2.

The observation equation can now be written as

$$\begin{pmatrix} Y(t) \\ Z(t) \end{pmatrix} = \begin{pmatrix} A \\ L \end{pmatrix} X(t) + \begin{pmatrix} u(t) \\ \epsilon(t) \end{pmatrix} \quad (3)$$

The joint distribution of $(X(1), X(2), \cdots)$ can be modeled through a state-space approach. Such an approach has been shown [15, 21] to perform well for modeling flow volumes. Specifically, we assume the following transition equation:

$$X(t) = CX(t-1) + w(t) \quad (4)$$

Further, let $Cov(w(t)) = \Sigma$. The state transition equation may suffer from mis-specification too and we look at this issue numerically in section 5.2. Note that the above equation corresponds to a vector auto-regressive model of order 1. Higher order models and model selection issues have been investigated by Zhao et. al. [26]. Flow modeling is an active area of research and one that is beyond the scope of this paper. Equation (12) is an extremely simple and robust way of modeling flow volumes that still captures spatio-temporal dependence which is key to filtering.

Finally, we will assume that there are certain budget and positivity constraints on the sampling rates $(\xi(1)', \xi(2)', \cdots)$. For the remainder of the paper, it is assumed that the constraints can be simplified as $\xi(t) \in \Xi(t)$, where $\Xi(t)$ are convex sets.

Given the state space model one can use a Kalman-Filter to recursively compute estimates of $X(t)$ based on information available at time t. Such an approach using only SNMP data $(Y(t)', Y(t-1)', \cdots)'$ has been been demonstrated in [21]. Let $\hat{X}(t)$ be the Kalman-Filter estimate of $X(t)$ based on $(Y(t)', Z(t)', Y(t-1)', Z(t-1)', \cdots)$ and let $\Sigma(t)$ denote the covariance matrix of estimation error i.e. $\Sigma(t) = E[(X(t)-\hat{X}(t))(X(t)-\hat{X}(t))']$.

The conditional covariance of $X(t)$ given $(Y(t-1)', Z(t-1)', Y(t-2)', Z(t-2)', \cdots)$ can be written as [11]:

$$\Sigma(t|t-1) = C\Sigma(t-1)C' + \Sigma \quad (5)$$

Further, we have the following updating equation:

$$\Sigma(t) = \Sigma(t|t-1) - \Sigma(t|t-1)J'F(t)^{-1}J\Sigma(t|t-1) \quad (6)$$

where $J' = (A', L')$ and

$$F(t) = J\Sigma(t|t-1)J' + \begin{pmatrix} \sigma^2 I & 0 \\ 0 & \Psi(t,\xi(t))^{-1} \end{pmatrix}$$

3. FORMULATION OF THE OPTIMAL SAMPLING PROBLEM

It is reasonable to seek to minimize some appropriate functional of $(\Sigma(1), \Sigma(2), \cdots)$, since this would lead to maximization of information about the underlying flows. More generally, if interest is restricted to a subset of flows $K'X(t)$, selected by a matrix K, then we seek to minimize a functional of $(K'\Sigma(1)K, K'\Sigma(2)K, \cdots)$. Notice that any *joint* minimization of $(K'\Sigma(1)K, K'\Sigma(2)K, \cdots)$ would have very high complexity, we restrict attention to optimizing $K'\Sigma(t)K$ at every time interval $t = 1, 2, \cdots$.

We will refer to the following problem as the optimal design problem for state space models. At each time interval t, the design objective would be to minimize $f(K'\Sigma(t)K)$ for some functional f. If $\mathbb{S}$ is the set of symmetric semi-definite matrices, then $f: \mathbb{S} \to \mathbb{R}$ is in general matrix isotonic [19]. The dimension of matrices in $\mathbb{S}$ is usually clear from the context and we will use the same notation $f()$ for all cases. Common choices of f are the determinant (also known as the D-optimality criterion in the design of experiments literature), the maximum eigenvalue (E-optimality) and the trace (A-optimality). Our exposition will focus on the E- and A- optimality criteria.

Remark: Although there is nothing special about these criteria conceptually, and in fact network engineers may find other criteria like relative mean squared error (MSE) attractive, we focus on these for the following reasons. First, these criteria result in optimization problems that are simple convex programs and thus intuitive and computationally easier to handle. Secondly, the purpose of this paper is to demonstrate that it is possible to get better performance with respect to specific design objectives though a targeted allocation of sampling resources. Generally any objective function would involve a scalarization of the covariance matrix of estimation error. Many reasonable scalarizations are well approximated by above criteria or simple modifications thereof. We discuss the connection between the optimality criteria and quantities like MSE in section 4.

Using the matrix inversion lemma ($(A+BDB')^{-1} = A^{-1} - A^{-1}B(B'A^{-1}B + D^{-1})^{-1}B'A^{-1}$), we can write

$$\Sigma(t)^{-1} = \Sigma(t|t-1)^{-1} + \sigma^{-2}A'A + L'\Psi(t, \xi(t))L \qquad (7)$$

The above can be interpreted as an information update equation [11]. Here $\Sigma(t)^{-1}$, the information available at time t, is represented as the sum of $\Sigma(t|t-1)^{-1}$, the information from time $t-1$, $\sigma^{-2}A'A$, information from SNMP data and $L'\Psi(t, \xi(t))L$ information from sampled data.

Now,

$$K'\Sigma(t)K = K'(\Sigma(t|t-1)^{-1} + \sigma^{-2}A'A + L'\Psi(t, \xi(t))L)^{-1}K \qquad (8)$$

Note that both $\Sigma(t)$ and $K'\Sigma(t)K$ are of the general form $G'M(\xi(t))^{-1}G$ where G is a known matrix and $M(\cdot)$ is a linear function. This will allow us to write the resulting optimizations as canonical convex programs.

A value of $\sigma = 0$, implies perfect knowledge in a specific subspace; i.e. $AX(t) = Y(t)$. We proceed by reparameterizing the variable $X(t)$ so that the equality $Y(t) = AX(t)$ can be solved for a subset of these new variables and the conditional covariance of the remaining variables is positive definite. Let $\tilde{X}(t) = Q'X(t)$, where Q is the unitary matrix from the eigenvalue decomposition

$$A'A = Q\Lambda Q'$$

and Λ being the diagonal eigenvalue matrix with first n_l diagonal values positive and the rest zero. The covariance of $\tilde{X}(t)$ given $(Y(t)', Z(t)', Y(t-1)', Z(t-1)', \cdots)$ can be written using the following proposition.

Proposition 1. *Let X and ϵ be independent normal random vectors with covariances Σ and Ψ^{-1} respectively. Further let $Y = AX$ for A with full row rank, $Z = LX + \epsilon$ and $A'A = Q\Lambda Q'$ be the eigenvalue decomposition of $A'A$. Let $Q = (Q_1, Q_2)$ such that (w.l.o.g) $AQ_1 \equiv P$, a full rank square matrix and $AQ_2 = 0$. Define*

$$\tilde{X} \equiv \begin{pmatrix} \tilde{X}_1 \\ \tilde{X}_2 \end{pmatrix} \equiv \begin{pmatrix} Q_1'X \\ Q_2'X \end{pmatrix}.$$

Then, the covariance of $\tilde{X}_1$ given Y is 0 and the covariance of $\tilde{X}_2$ given Y and Z is $(Q_2'\Sigma^{-1}Q_2 + Q_2'L'\Psi LQ_2)^{-1}$.

Proof. Note that $AX = P\tilde{X}_1$ and thus $\tilde{X}_1 = P^{-1}Y$. Hence, clearly the covariance of $\tilde{X}_1$ given Y is 0. Further, the covariance of $\tilde{X}_2$ given $(Y', Z')'$ is the same as $\Sigma_{\tilde{X}_2|(\tilde{X}_1', \tilde{Z}')'}$, the covariance of $\tilde{X}_2$ given $(\tilde{X}_1', Z')'$. Let $\Sigma_{\tilde{X}_2|\tilde{X}_1}$, $\Sigma_{\tilde{X}_2,Z|\tilde{X}_1}$ and $\Sigma_{Z|\tilde{X}_1}$ be the covariance of $\tilde{X}_2$, cross-covariance of $\tilde{X}_2$ and Z and the covariance of Z respectively given $\tilde{X}_1$. Then

$$\begin{aligned}\Sigma_{\tilde{X}_2|(\tilde{X}_1', \tilde{Z}')'} &= \Sigma_{\tilde{X}_2|\tilde{X}_1} - \Sigma_{\tilde{X}_2,Z|\tilde{X}_1}\Sigma^{-1}_{Z|\tilde{X}_1}\Sigma'_{\tilde{X}_2,Z|\tilde{X}_1} \\ &= \Sigma_{\tilde{X}_2|\tilde{X}_1} - \\ \Sigma_{\tilde{X}_2|\tilde{X}_1}Q_2'L'(LQ_2\Sigma_{\tilde{X}_2|\tilde{X}_1}&Q_2'L' + \Psi^{-1})^{-1}LQ_2\Sigma_{\tilde{X}_2|\tilde{X}_1} \\ &= (\Sigma^{-1}_{\tilde{X}_2|\tilde{X}_1} + Q_2'L'\Psi LQ_2)^{-1} \end{aligned} \qquad (9)$$

The last equality follows from the matrix inversion lemma. Further

$$\Sigma_{\tilde{X}_2|\tilde{X}_1} = Q_2'\Sigma Q_2 - Q_2'\Sigma Q_1(Q_1'\Sigma Q_1)^{-1}Q_1'\Sigma Q_2$$

Therefore,

$$\begin{aligned}\Sigma^{-1}_{\tilde{X}_2|\tilde{X}_1} &= [(Q'\Sigma Q)^{-1}]_{22} \\ &= [Q'\Sigma^{-1}Q]_{22} = Q_2'\Sigma^{-1}Q_2\end{aligned}$$

which establishes the desired result. □

Hence the covariance of $\tilde{X}(t)$ given $(Y(t)', Z(t)', Y(t-1)', Z(t-1)', \cdots)$ is given by

$$\tilde{\Sigma}(t) = \begin{pmatrix} 0 & 0 \\ 0 & (Q_2'\Sigma(t|t-1)^{-1}Q_2 + Q_2'L'\Psi(t, \xi(t))LQ_2)^{-1} \end{pmatrix}$$

Finally,

$$\begin{aligned}\Sigma(t) = Q\tilde{\Sigma}(t)Q' &= Q_2(Q_2'\Sigma(t|t-1)^{-1}Q_2 \\ &+ Q_2'L'\Psi(t, \xi(t))LQ_2)^{-1}Q_2'\end{aligned} \qquad (10)$$

which shows that both $\Sigma(t)$ and $K'\Sigma(t)K$ are of the form $G'M(\xi(t))^{-1}G$.

4. OPTIMALITY CRITERIA AND CONVEX PROGRAMS

For the classical optimality criteria mentioned earlier, the optimization problems can be written as canonical convex programs [14] that can be solved using any of the commonly available software packages. The usual class of algorithms used for solving convex optimization problems, Interior Point Methods, have polynomial time complexity and have been successfully employed for various large scale optimization problems. Since we focus on one measurement interval at a time, the general optimization problem is to minimize $f(G'M(\xi)^{-1}G)$ subject to $\xi \in \Xi$.

4.1 E-optimality

E-optimality corresponds to minimizing the maximum eigenvalue of a covariance matrix. Equivalently, this minimizes the worst possible variance of any linear combination of the error vector. Roughly, this would lead to a small value for the largest of the MSE.

In this case we get the following program:

$$\underset{\theta \in \mathbb{R}, \xi \in \Xi}{\text{maximize}}\, \theta$$

subject to

$$G'M(\xi)^{-1}G \preceq \theta^{-1}I \qquad (11)$$

Recall that if $A \succ 0$ and $C \succ 0$ then

$$\begin{pmatrix} A & B \\ B' & C \end{pmatrix} \succeq 0$$

if and only if $C \succeq B'A^{-1}B$ or equivalently if and only if $A \succeq BC^{-1}B'$. This follows from taking the Schur complement of C and A respectively.

Thus, the constraint (11) is equivalent to

$$M(\xi) \succeq \theta GG'$$

The above program is a Semi-Definite Program (SDP) if Ξ is a convex polygon [2].

4.2 A-optimality

A-optimality corresponds to minimizing the trace of a covariance matrix and hence minimizing the average MSE. In this case we get the following program:

$$\underset{t \in \mathbb{R}^n, \xi \in \Xi}{\text{minimize}} \sum_{i=1}^{n} t_i$$

subject to

$$t_i \geq e_i' G' M(\xi)^{-1} G e_i$$

for $i = 1, \cdots, n$, where n is the number of columns in G and e_i is the ith unit-vector of $\mathbb{R}^n$. Using Schur complement we can write the given constraints as:

$$\begin{pmatrix} M(\xi) & Ge_i \\ e_i'G' & t_i \end{pmatrix} \succeq 0$$

Thus we get a SDP if Ξ is a convex polygon.

4.3 D-optimality

D-optimality corresponds to minimizing the determinant of a covariance matrix and hence minimizing the volume of the associated confidence ellipsoid. In this case we get the following program:

$$\underset{\xi \in \Xi}{\text{minimize}} \log \det G' M(\xi)^{-1} G$$

Note that if G does not have full column rank then the minimization is trivially unbounded. Although D-optimality has a natural interpretation for polynomial regression models we will focus on A- and E- optimality due to their simplicity and interpretability in terms of error variances.

5. PERFORMANCE EVALUATION

In this Section, the performance of the proposed design scheme for the sampling rates is evaluated. Specifically, the following issues are addressed: in the first part, the cumulative effect of the greedy design optimization steps is examined, when the true model of flow volumes is given by the state-transition equation (4) and errors in the observation equation (3) have constant variance. In the second part, the effects of model misspecification in the state transition equation and approximations in the supposed variances of errors of the observation equation are investigated.

In the following, we describe the overall setup for the numerical study.

1. *Topology, Routing and Observation Points.* The Abilene network topology (figure 2) is used in our experimental setup. It consists of 11 nodes and $16 \times 2 = 32$ directed edges between pairs of nodes (bidirectional links). Flows exist between all pairs of nodes resulting in a total of $11 \times 10 = 110$ flows. We assume that these flows are routed through minimum distance paths. Further we assume that SNMP data are available from all the edges. Similarly, sampled data are collected at each edge. All the incoming edges at a node are considered as the interfaces of the corresponding router.

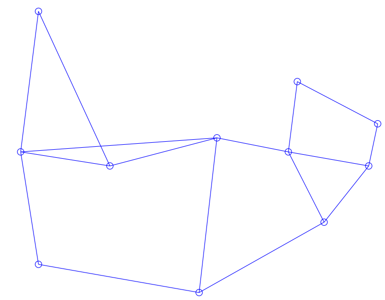

Figure 2: Abilene Topology used for Performance Evaluation Purposes

2. *Budget Constraints.* The budget constraint represents an effective way of limiting the number of samples collected at a particular router. For this numerical study, we use a budget constraint of the form $R\xi_t \leq b$. Here R_{ij} is equal to the square of the number of flows traversing edge j if edge j is incident on node i and 0 otherwise. Thus, the weighted sum of sampling rates on the incoming edges of node i is bounded above by b_i. It can be seen that the cost of sampling on an edge goes up as the square of number of flows traversing that edge. This is a very reasonable assumption, because not only does sampling on a heavily used link results in a large number of samples and thus a significant demand on resources, but also on such a link less resources are available for sampling to begin with. The elements of b were identically chosen to be some b_0 to ensure realistic sampling rates. The above router level constraint implies the network level constraint $\mathbf{1}'R\xi_t \leq \mathbf{1}'b$. While most of the following assumes router level constraints we will also investigate the result of imposing only the network level budget constraint.

Notice that the above setup has the following symmetry property. Define a short flow to be a flow that traverses exactly one edge or equivalently exactly one router interface. Further, each interface is traversed by exactly one short flow. By definition, any given interface is the only point in the network from where sampled data is available for the short flow traversing it.

The above symmetry has important consequences for the optimization problems described in the previous section. An intuitive understanding of its effects can be described as follows. A low sampling rate on any interface would result in an information deficit on the corresponding short flow, which in turn will result in a large measurement error variance due to the inverse relation between variance and sampling rate. From equation (7), the total information at time t is the sum of information obtained from the previous time period $t-1$, from SNMP measurements and from sampled data. Suppose that the sum of information from time $t-1$ and from SNMP measurements do not ameliorate the problem of information deficit on short flows. Further, suppose that there is roughly equal information from this sum on each short flow. In this case, the objective function of A-optimality (i.e. average estimation error variance) is heavily influenced by the large measurement variances of short flows. Further, the objective function of E-optimality is determined to an even larger extent by the largest measurement error variance. Thus, to the extent possible given the budget constraints, all the short flows must be sampled roughly equally under the optimal allocation. This in turn implies that all interfaces should be allocated roughly equal sampling rates, again to the extent possible under the budget constraints. This was indeed found to be true as described in the following.

On the other hand, when interest focuses on a specific subsystem of flows $K'X(t)$, then it is not a priori clear what an optimal or near optimal solution should be. For most of following, we assume that the interest is restricted to long flows; i.e. flows that traverse four or more edges in the employed topology. A matrix K was used to subset $X(t)$ accordingly.

5.1 Cumulative Effect of Greedy Optimality

Recall that at each time interval t, the measurement scheme was designed to better utilize the side information available from past measurements $\Sigma(t|t-1)^{-1}$ and from SNMP measurements $\sigma^{-2}A'A$. The gains compared to a homogeneous allocation may be small in a single step but these gains propagate across time in the form of prior information according to equation (5). The purpose of this section is to study how the gains from using an optimal design accumulate over time.

When the assumed observation and state transition models (34) are the true model and the true C, Σ, σ and $\Psi(t,\cdot)$ matrices are known, then $\Sigma(t)$ calculated using (56) is the true covariance of estimation error. Hence, we can calculate the covariance of estimation error for any sequence of allocations $(\xi(1)', \xi(2)', \cdots)$.

For this investigation we assume $\sigma = 0$, $C = I$, $\Sigma = I$ and $\Psi(t,\xi) = \Psi(\xi) = \sum_k \xi_k \Psi_k$. The important calibration issue here is to choose Ψ_k to maintain realistic order of magnitude of measurement error $\Psi(\xi)^{-1}$ with respect to the innovations covariance Σ. Empirical data has shown [15] that the innovations in flow volumes have a variance of roughly the same magnitude as the flow volume itself. Therefore, it is expected that $[\Sigma]_{ii}$ is of the same order of magnitude as X_i. Using this observation and equation (12) we choose

$$[\Psi_k]_{ii} = \begin{cases} 1 & \text{if } i \in \mathcal{C}_k \\ 0 & \text{otherwise} \end{cases}$$

For the sake of comparison, a set of naive allocations is defined as follows. If we are interested in the estimation error for all flows and there is a router level budget constraint, it is given by ξ for which each interface on a router has equal sampling rate and $R\xi = b$. Note that interfaces on different routers can have different sampling rates. This ξ is also used as the common initial allocation $\xi(0)$ for all allocation schemes considered in the paper. On the other hand, if there is only a network level budget constraint and we are interested in the estimation error for all flows, the naive allocation correspond to ξ for which all interfaces in the network have equal sampling rates and $\mathbf{1}'R\xi = \mathbf{1}'b$. Finally,

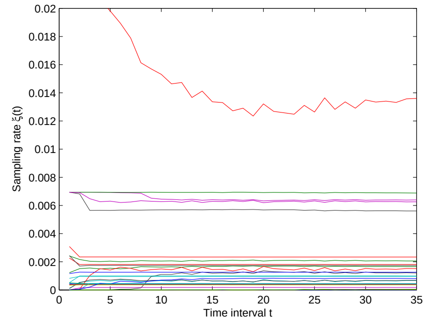

Figure 4: Evolution of E-optimal sampling rates

when we are interested in the estimation error in $K'X(t)$ (the subset of long flows) the naive allocation is given by ξ for which any interface that is not traversed by any long flow has 0 sampling rate, all other interfaces have sampling rates determined as above for router level and network level budget constraints, respectively.

We consider two sets of initial conditions. The first one, referred to as noise-free initialization, assumes perfect knowledge at time 0 ($\Sigma(0) = 0$). The second, referred to as noisy initialization, assumes

$$\Sigma(0) = Q_2(Q_2'L'\Psi(\xi(0))LQ_2)^{-1}Q_2'$$

From equation (10), the latter corresponds to starting with only one measurement interval worth of information. All results are for router level budget constraints, unless otherwise stated.

First consider the case when we are interested in estimation errors of all flows. Figures 3 (a) and (b) show the value of the objective function for E-optimality, i.e. the largest eigenvalue of $\Sigma(t)$, over time for 3 Kalman filters. These filters respectively use SNMP data only, SNMP and naively sampled data and SNMP and optimally sampled data. Clearly inclusion of sampled data gives substantial improvement over SNMP alone. Moreover, as information accumulates over time we get an improvement in performance. This can be viewed as the gain from using a temporal model to combine samples from across time intervals. For the E-optimality criterion, the performance of the naive and optimal allocations is fairly similar. This is not spritzing given the symmetry in our setup, discussed above. Figure 4 presents the evolution of the allocations $\xi(t)$ over time for noisy initialization. Clearly the sampling rates achieve a steady state for most of the interfaces after very few time intervals (3-5) and certainly before $t \approx 15$.

The above qualitative assessments hold for both the E- and A- optimality criteria. As an example, Figure 5 shows the value of the objective function for A-optimality, i.e. the trace of $\Sigma(t)$ over time starting from noisy initialization based on the three Kalman filters employed. In this case, the third filter uses A-optimally sampled data. Notice that the Kalman filter based on optimally sampled data has slightly better transient behavior than the one based on naively sampled data. The difference from the E-optimal case (figure 3(a)) is due to the fact that A-optimality is less sensitive to high variance or low information on particular flows. Hence, the A-optimal solution is driven by short flows to a lesser degree than the E-optimal solution, despite the

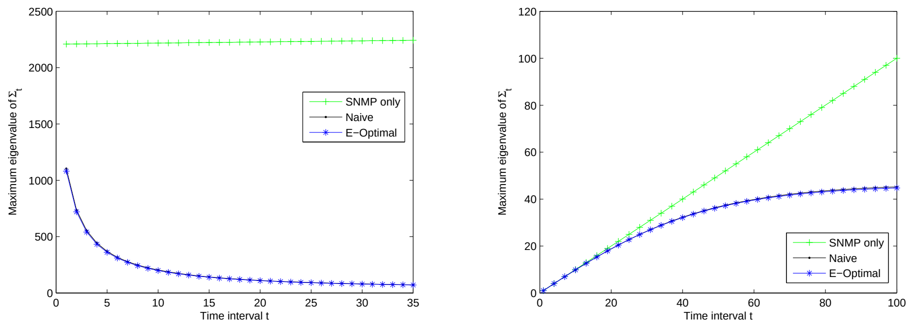

(a) **(b)**

Figure 3: Performance of E-optimal sampling for (a)noisy and (b) noise-free initialization

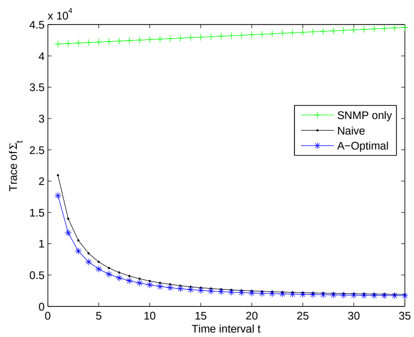

Figure 5: Performance of A-optimal sampling for noisy initialization

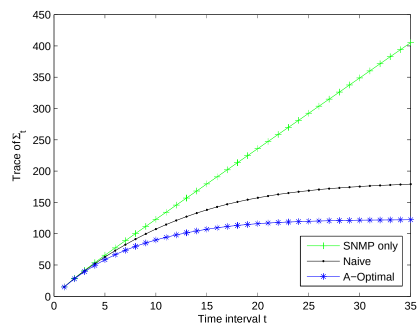

Figure 8: Performance of A-optimal sampling under network level constraint

symmetry property mentioned earlier.

Next, we focus on the case when the interest is on $K'X(t)$, the subset of long flows. Figures 6 (a), (b), (c) and (d) show the relevant performance metric over time for E- and A- optimal sampling strategies for the two sets of initial conditions. In this case, we see a clear advantage in using optimal sampling over naive sampling. In respective performance metrics the improvement over naive sampling is 29.7% and 26.0% for E-optimal design and 18.0% and 15.7% for A-optimal design.

Next, we look more closely at the steady state sampling rates. Figures 7(a) and (b) show the evolution and spatial profile of the final E-optimal sampling rates under noisy initialization. In Figure 7 (c) the ξ_k values versus k for the above allocation and the naive allocation are shown. Clearly the optimal allocation has a sparser support. Intuitively all interfaces that have a high cost but do not provide enough information on the *long flows* have zero sampling rate. Figure 7 (d) shows for each flow the cumulative sampling rate over all observation points that the flow traverses for naive and optimal allocation, respectively. Notice that the minimum allocation to a long flow is larger under optimal sampling (.0046) than for naive sampling (.0033). It clearly illustrates the benefits obtained from being able to design the sampling rates in an optimal manner, which can be responsive to the objective set forth by network engineers.

Finally, we look at some variations on the above setup. Figure 8 shows the performance of A-optimal sampling from a noise-free initialization under a *network-level* constraint alone. Here the improvement over naive sampling is 31.8%. Clearly there is improvement in performance over the corresponding performance under router-level constraints (figure 6(d)). Another variation is to replace the matrix R by its element-wise square-root. Thus, the cost of sampling on an interface goes up linearly with the number of flows traversing that interface. Notice that from a geometric point of view we have not changed the nature of the constraints. Figure 9 shows the performance of E-optimal sampling under such a constraint. Notice that compared to our usual cost (figure 6(a)) the gain over the corresponding naive allocation is smaller. This is perhaps not surprising since under linear cost the gains from sampling on an interface are roughly equal to the cost and hence any shift along the surface $\{\xi : R\xi = b\}$ would have small effect on the objective function.

5.2 Effects of Misspecification

The state transition equation (4) and observation equation (3) necessarily suffer from model misspecification and parameter inaccuracies. Note that if $X(t)$ has a constant covariance, then (4) can be true only if the spectral radius of C is less than 1 [16]. However, in this case, the steady state mean of $X(t)$ would be 0. This model (4) works well in practice [21] possibly because the misspecification in the conditional distribution of $X(t)$ given $X(t-1)$ is small. The misspecification in the observation equation (3) may potentially be more serious. Firstly, the sampled data $Z(t)$ are

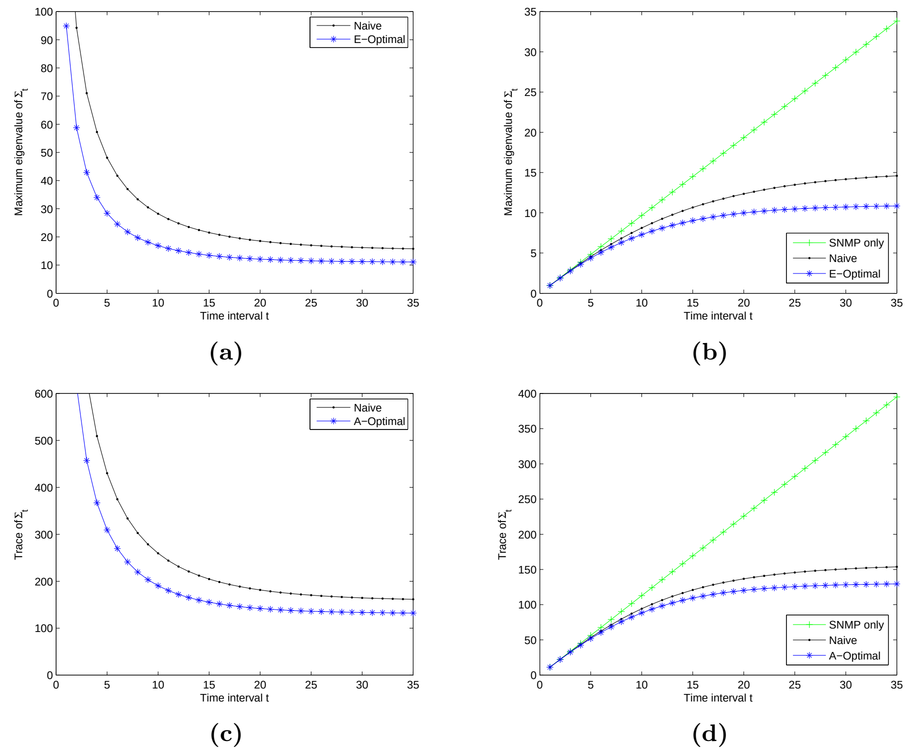

Figure 6: Performance of E-optimal (a)(b) and A- optimal sampling for noisy (a)(c) and noise-free (b)(d)initialization

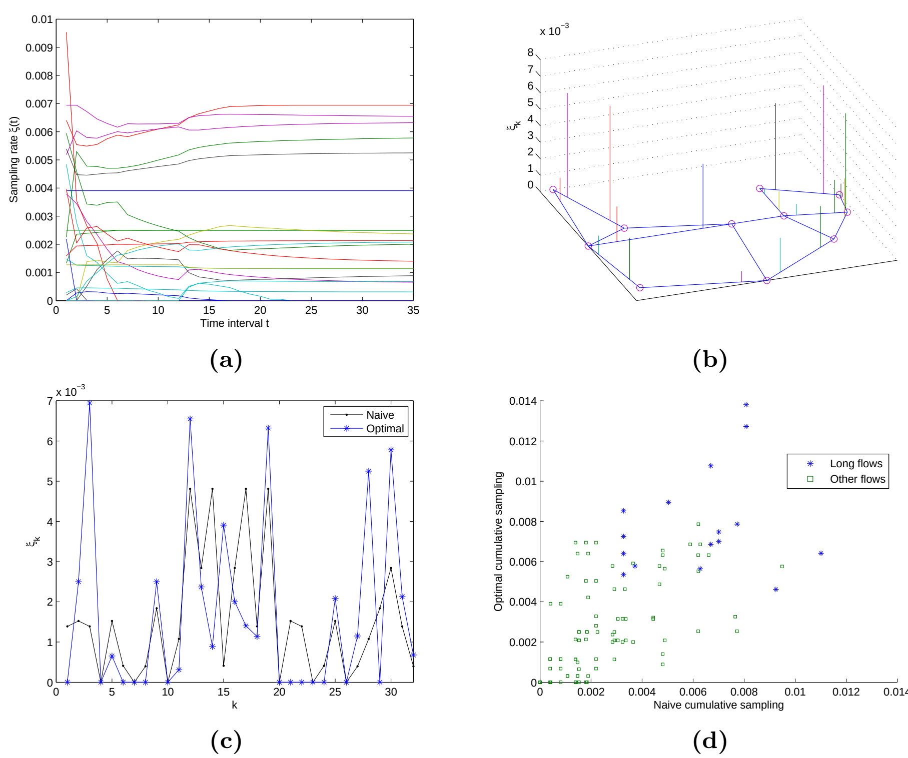

Figure 7: Spatio-temporal behavior of E-optimal ξ

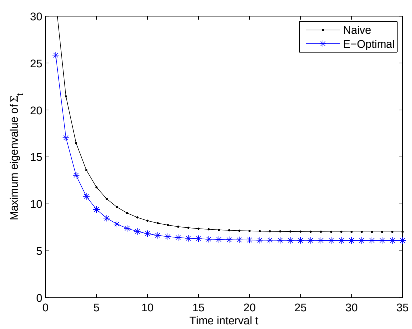

Figure 9: Performance of E-optimal sampling under linear cost.

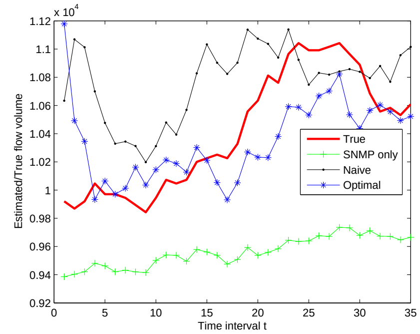

Figure 10: Sample trajectory of true and estimated flow volume

not distributed normally given $X(t)$ to begin with. Thus, Kalman filtering is not optimal. Second, it is clear from equation (12) that the true value of the covariance of the measurement error, $\Psi(t,\cdot)$, can not be known with high precision and one would use an approximation for both filtering and design optimization. We study the effects of such misspecification in the following.

We start by describing the model used for data generation. Assume,

$$X_i(t) = \text{ROUND}(aV_i(t) + \cdots + aV_i(t+n))$$

where $V_i(t)$ are independent Poisson random variables with mean λ_i. Variables a and λ_i are chosen to ensure $E[X_i(t)] = E[(X_i(t) - X_i(t-1))^2] = \mu_i$, where μ_i is the required mean. Thus, we again ensure that the variance of innovations for a flow are of the same magnitude as the volume of the flow, as seen in empirical studies [15]. The above can be solved for $a = (n+1)/2$ and $\lambda_i = 2\mu_i/(n+1)^2$ Further, the above model implies that the corresponding fitted model for the state transition equation (4) would have $C = I$ and $\Sigma = \text{diag}(\mu)$. These are the values we use for both design optimization and filtering. For most of the results presented here it assumed that $n = 50$, while the μ_i's were chosen from a uniform distribution with support between 9000 and 11000.

The sampled measurements were generated as follows. The SNMP measurements $Y(t)$ were generated from a normal distribution with mean $X(t)$ and covariance $\sigma^2 I$ with $\sigma^2 = 100$. The i-th sampled measurement, $Z_i(t)$ was chosen to be $1/\xi_k(t)$ times a binomial random variable with parameters $X_j(t)$ (the corresponding flow volume) and $\xi_k(t)$ (the corresponding sampling rate), where $j = l(i)$ and k is such that $\mathcal{C}_k \ni i$ (see section 2).

Finally, for the purpose of design optimization and filtering it is assumed that $\psi(t,\xi) = \sum_k \xi_k \Psi_k$, where Ψ_k are diagonal matrices with

$$[\Psi_k]_{ii} = \begin{cases} 1/\hat{X}_{l(i)}(t-1) & \text{if } i \in \mathcal{C}_k \\ 0 & \text{otherwise} \end{cases} \qquad (12)$$

where $\hat{X}(t-1)$ is the Kalman filter estimate of $X(t-1)$. This choice works well if $\hat{X}(t-1)$ is a good estimate of $X(t)$. However, if $\hat{X}(t-1)$ is based on a few intervals worth of data, as in the first few time intervals after a noisy initialization, it can be extremely noisy. In our simulations rare instances of explosive growth in estimation error were observed[1].

[1]These instances were extremely infrequent, occurring at most thrice in 200×2 realizations of sampled data (corresponding to optimal and naive sampling of the 200 realizations of $X(t)$). We have dropped such cases from our results and indicated accordingly

In this setup, two sets of initial conditions were considered as well. The noise-free initialization assumes $\hat{X}(0) = X(0)$ and $\Sigma(0) = 0$, while the noisy one assumes

$$\hat{X}(0) = (J'J)^{-1}J' \begin{pmatrix} Y(0) \\ Z(0) \end{pmatrix}$$

where $J' = (A', L')$ and

$$\Sigma(0) = (J'J)^{-1}J' \begin{pmatrix} \sigma^2 I & 0 \\ 0 & \Psi^{-1}(0,\xi(0)) \end{pmatrix} J(J'J)^{-1}$$

We define $\hat{X}(-1) \equiv \hat{X}(0)$ for the purpose of evaluating $\Psi(0,\cdot)$. Also, recall that $\xi(0)$ is the naive allocation, which was used to generate $Z(0)$.

For a given optimization setup we generate μ as described and simulated up to 200 realizations of $X(t)$. We calculate $\hat{X}(t)$ for different versions of the Kalman filter operating on different data depending on the respective sampling rates. Figure 10 depicts one realization of the trajectories of the true flow volume and estimated flow volumes for a particular flow. Thus, we calculate the MSE for all flows, averaging over all realizations. Given the discussion in section 4, we study the performance of the maximum MSE among all flows for E-optimality and the mean of MSEs of all flows for A-optimality.

Figure 11 depicts the performance of optimal sampling for different initializations and different target sets of flows (for (a) and (c) two and one simulation, respectively, out of 200 were dropped due to numerical instability). Comparing Figures 11(a)(b) to Figures 3 (a)(b), and Figures 11(c)(d) to figures 6 (a)(b) we observe that the relative performance of different Kalman filters under misspecification is very similar to that of the corresponding set of Kalman Filters when there is no misspecification. Specifically, we again observe that when interest is restricted to the subset of long flows there is optimal sampling performs considerably better than naive sampling.

Finally, we look at some variations on the above settings. Recall that $\Sigma = \text{diag}(\mu)$ and thus one needs knowledge of μ for both filtering and tracking purposes. We study the effect of incorrectly specified μ on the performance of the Kalman filter. Figure 12 shows the performance of re-

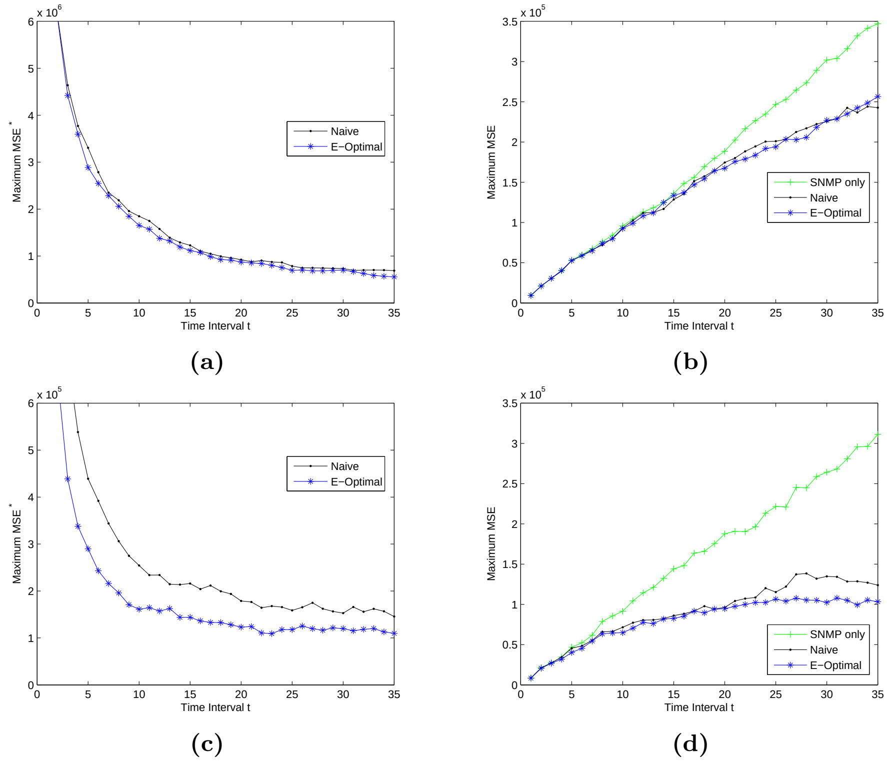

Figure 11: Performance of E-optimal for noisy (a)(c) and noise-free (b)(d) realizations. Interest is restricted to long flows in (c)(d).

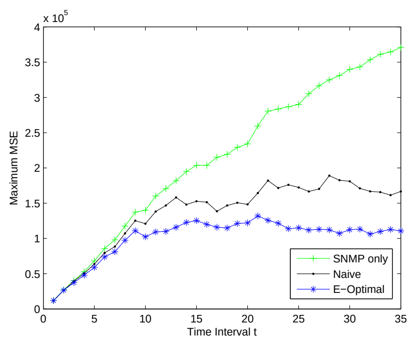

Figure 12: Performance under misspecification of μ

spective Kalman filters when the presumed μ is still chosen as described earlier, but the true μ used for data generation has two arbitrary entries changed to 20000. Notice that for the above misspecification the relative improvement from optimal sampling over naive sampling is considerably larger than the corresponding improvement otherwise (Figure 11(d)). Figure 13 shows the performance of E-optimal sampling when R is replaced by its element-wise square-root. Given our observation from figure 9 it is not surprising that the relative gain from optimal sampling over naive sampling is reduced.

6. CONCLUSIONS AND FUTURE WORK

In this paper, the problem of tracking network flow volumes over time has been addressed in the context of a state-space model, which is of great importance in network monitoring applications. It was shown that collecting direct measurements on the flows through sampling leads to significant reductions in the estimation error. Further, these gains from sampling can be maximized by choosing the sampling rates in an optimal manner. A significant theoretical contribution of this study is the development of the theoretical framework for optimally designing the sampling rates. This corresponds to an optimal design of the observation equation. Clearly the gains from step-wise optimality add up to give significant improvements. The extensive numerical study undertaken quantifies the obtained gains in performance. We find that the qualitative conclusions for simple linear models continue

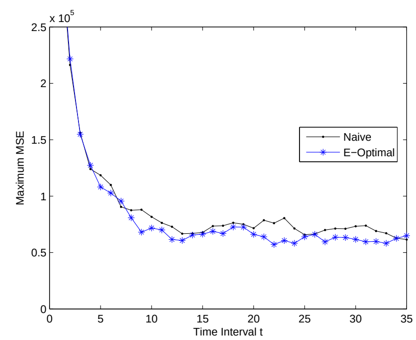

Figure 13: Performance of E-optimality under linear cost

to hold for simulated data that closely resembles network data.

A number of interesting problems remains open and require further study. The linear state space model is fairly analytically tractable and one would like to study more closely the evolution of optimal sampling rates. Specifically we would like to investigate the case when sampling rates cannot be changed during every measurement interval and are instead constant over blocks. A myopic approach can in principle still be used with some deterioration in performance. In general, network monitoring involves quantities for which a linear filter represents a convenient mathematical approximation. It is of great theoretical and practical interest to study the problem of designing optimal sampling rates in the context of non-linear filtering and also involving non-linear cost and objective functions.

7. REFERENCES

[1] P. Barford, J. Kline, D. Plonka, and A. Ron. A signal analysis of network traffic anomalies. In *IMW '02: Proceedings of the 2nd ACM SIGCOMM Workshop on Internet measurment*, pages 71–82, New York, NY, USA, 2002. ACM Press.

[2] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.

[3] J. Cao, D. Davis, S. Wiel, and B. Yu. Time-varying network tomography : Router link data. *Journal of the American Statistical Association*, 95:1063–75, 2000.

[4] A. Chen and J. Cao. Method of one dimensional projections for network tomography. *Statistical Inverse Problems, R. Liu, W. Strawderman, C. H. Zhang (editors), IMS Lecture Note Series*, 2006.

[5] B. Choi and S. Bhattacharrya. On the accuracy and overhead of Cisco sampled netflow. In *ACM Sigmetrics Workshop on Large-Scale Network Inference (LSNI)*, Banff, Canada, June 2005.

[6] N. Duffield, C. Lund, and M. Thorup. Properties and prediction of flow statistics from sampled packet streams. In *IMW '02: Proceedings of the 2nd ACM SIGCOMM Workshop on Internet measurment*, pages 159–171, New York, NY, USA, 2002. ACM Press.

[7] N. Duffield, C. Lund, and M. Thorup. Flow sampling under hard resource constraints. In *SIGMETRICS '04/Performance '04: Proceedings of the joint international conference on Measurement and modeling of computer systems*, pages 85–96, New York, NY, USA, 2004. ACM.

[8] N. Duffield, C. Lund, and M. Thorup. Optimal combination of sampled network measurements. In *IMC'05: Proceedings of the Internet Measurement Conference 2005*, pages 8–8, Berkeley, CA, USA, 2005. USENIX Association.

[9] N. G. Duffield, C. Lund, and M. Thorup. Learn more, sample less: control of volume and variance in network measurement. *IEEE Transactions on Information Theory*, 51(5):1756–1775, 2005.

[10] V. Erramilli, M. Crovella, and N. Taft. An independent-connection model for traffic matrices. In *IMC '06: Proceedings of the 6th ACM SIGCOMM on Internet measurement*, pages 251–256, New York, NY, USA, 2006. ACM Press.

[11] A. C. Harvey. *Forecasting, structural time series models and the Kalman filter*. Cambridge University Press, 1990.

[12] A. Lakhina, M. Crovella, and C. Diot. Diagnosing network-wide traffic anomalies. In *SIGCOMM '04: Proceedings of the 2004 conference on Applications, technologies, architectures, and protocols for computer communications*, pages 219–230, New York, NY, USA, 2004. ACM Press.

[13] E. Lawrence, G. Michailidis, V. Nair, and B. Xi. Network tomography: a review and recent developments. In *Frontiers in Statistics*, pages 345–364. Fan and H. Koul (eds), Imperial College Press, London, 2006.

[14] J. Lee. Semidefinite programming in experimental design. In *H. Wolkowicz, R. Saigal and L. Vandenberghe, editors, "Handbook of Semidefinite Programming", International Series in Operations Research and Management Science, Vol. 27*. Kluwer, 2000.

[15] G. Liang, N. Taft, and B. Yu. A fast lightweight approach to origin-destination IP traffic estimation using partial measurements. *IEEE Transactions on Information Theory*, 52(6):2634–2648, 2006.

[16] H. Lutkepohl. *Introduction to multiple time series analysis*. Springer-Verlag, New York, 1993.

[17] A. Medina, N. Taft, K. Salamatian, S. Bhattacharyya, and C. Diot. Traffic matrix estimation: existing techniques and new directions. In *Proceedings of the 2002 conference on Applications, technologies, architectures, and protocols for computer communications*. ACM Press, 2002.

[18] L. Peterson and B. Davie. *Computer Networks: A Systems Approach*. Morgan Kaufmann, San Francisco, 2003.

[19] F. Pukelsheim. *Optimal Design of Experiments*. John Wiley & Sons, Inc., 1993.

[20] H. Singhal and G. Michailidis. Identifiability of flow distributions from link measurements with applications to computer networks. *Inverse Problems*, 23(5):1821–1849, 2007.

[21] A. Soule, A. Lakhina, N. Taft, K. Papagiannaki, K. Salamatian, A. Nucci, M. Crovella, and C. Diot. Traffic matrices: Balancing measurements, inference and modeling. In *Proceedings of the joint international conference on Measurement and modeling of computer systems*. ACM Press, 2005.

[22] A. Soule, A. Nucci, R. Cruz, E. Leonardi, and N. Taft. How to identify and estimate the largest traffic matrix elements in a dynamic environment. In *SIGMETRICS '04/Performance '04: Proceedings of the joint international conference on Measurement and modeling of computer systems*, pages 73–84, New York, NY, USA, 2004. ACM Press.

[23] L. Yang and G. Michailidis. Estimation of flow lengths from sampled data. In *Proceedings of IEEE GLOBECOM*, 2006.

[24] L. Yang and G. Michailidis. Sample based estimation of network traffic characteristics. In *Proceedings of IEEE INFOCOM*, 2007.

[25] Y. Zhang, M. Roughan, N. G. Duffield, and A. G. Greenberg. Fast accurate computation of large-scale IP traffic matrices from link loads. In *Proceedings of*

the International Conference on Measurements and Modeling of Computer Systems, SIGMETRICS 2003, pages 206–217, 2003.

[26] Q. Zhao. *Towards Ideal Network Traffic Measurement: A Statistical Algorithmic Approach*. PhD thesis, College of Computing, Georgia Institute of Technology, Atlanta, Georgia, USA, 2007.

[27] Q. Zhao, Z. Ge, J. Wang, and J. Xu. Robust traffic matrix estimation with imperfect information: making use of multiple data sources. In *SIGMETRICS '06/Performance '06: Proceedings of the joint international conference on Measurement and modeling of computer systems*, pages 133–144, New York, NY, USA, 2006. ACM.