OPTIMAL SAMPLING STRATEGIES FOR MULTISCALE MODELS WITH APPLICATION TO NETWORK TRAFFIC ESTIMATION

Vinay J. Ribeiro, Rudolf H. Riedi and Richard G. Baraniuk

Department of Electrical and Computer Engineering, Rice University
6100 South Main Street, Houston, TX 77005, USA

ABSTRACT

This paper considers the problem of determining which set of 2^p leaf nodes on a binary multiscale tree model of depth N ($N > p$) gives the best linear minimum mean-squared estimator of the tree root. We find that the best-case and worst-case sampling choices depend on the correlation structure of the tree. This problem arises in Internet traffic estimation, where the goal is to estimate the average traffic rate on a network path based on a limited number of traffic samples.

1. INTRODUCTION

Computer networks, due to their sheer size and complexity, give rise to several challenging signal processing problems. Networks like the Internet transfer data (or traffic) in the form of packets from one point to another through routers (see Fig. 1). Because privacy and security issues prevent the public sharing of measurements made at routers, network users are forced to solve inverse problems to infer internal network properties via packet delay measurements made solely at the edge of the network.

An important inverse inference problem we study is that of *traffic estimation*. Consider the simple network path consisting of a traffic process $X(t)$ entering a single router queue with service rate C bits/s depicted in Fig. 2. The traffic X is typically a bursty random process possessing fractal properties like self-similarity and long-range-dependence (LRD) [1].

Our aim is to estimate $R = \int_0^T X(t)dt$, the total traffic volume entering the queue in time interval $[0, T]$, from the spacings of specially injected packets called *probes*. If we inject two probe packets of size P bits into the queue at times $t = \tau$ and $t = \tau + \Delta$ where $\Delta = P/C$, also known as a *packet-pair*, then their spacing after the queue directly reflects the amount of traffic arriving between them, that is $\int_\tau^{\tau+\Delta} X(t)dt$. Obviously one can estimate R by saturating the queue with packet-pairs for the entire time interval $[0, T]$. This solution is highly impractical since we would use up the entire bandwidth with probe traffic and congest the network! We are thus forced to estimate R using only a few probe packet-pairs. Since X is not bandlimited in general, the problem of estimating R from few traffic samples is ill-posed.

This work was supported by NSF grant ANI–0099148, DoE SciDAC grant DE-FC02-01ER25462, DARPA/AFRL grant F30602–00–2–0557, and the Texas Instruments Leadership University program.
Email: {vinay, riedi, richb}@rice.edu, URL: dsp.rice.edu.

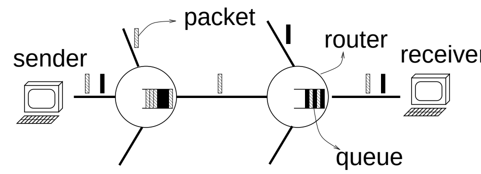

Fig. 1. *The Internet transmits data packets from sender to receiver.*

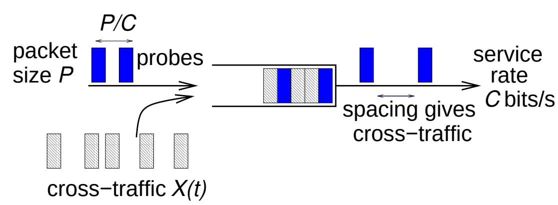

Fig. 2. *Probing for competing traffic. The spacing between probe packets gives the volume of competing traffic arriving between them.*

In this paper we determine what spacing between different probe packet-pairs gives the optimal linear minimum mean squared estimate (LMMSE) of R. Should packet-pairs be spaced uniformly apart in $[0, T]$ or bunched together at one end of $[0, T]$? We restrict our analysis to traffic well-modeled by a multiscale tree. In related work, the authors of [2] numerically compared different probing schemes for certain stationary traffic models.

The paper is organized as follows. We describe multiscale tree models in Section 2 and state our optimality results and a conjecture in Section 3. After presenting numerical simulations in support of the conjecture in Section 4, we conclude in Section 5. The Appendix contains the proofs.

2. MULTISCALE TREE MODELS

We model the traffic process X using a multiscale dyadic tree (see Fig. 3). The tree nodes are defined by

$$V_{j,k} := \int_{Tk2^{-j}}^{T(k+1)2^{-j}} X(t)dt \qquad (1)$$

for $j = 0, \ldots, N;\ k = 0, \ldots, 2^j - 1$. The tree-root value $V_{0,0}$ thus equals R. In this paper we focus on 2^{nd}-*order statistical trees*.

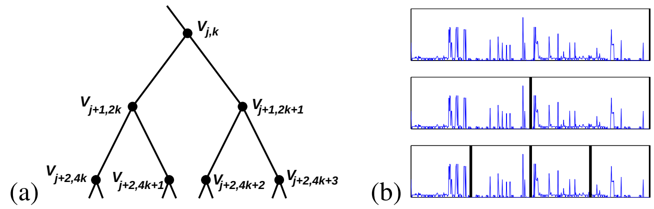

Fig. 3. *Multiscale tree representation of a traffic trace. Nodes at each horizontal level in the tree in (a) correspond to the sum (aggregates) of the process X in non-overlapping blocks of sizes of powers of two as in (b), with lower levels corresponding to smaller block sizes. Each node is the sum of its two child nodes.*

Definition 1 *V is a 2^{nd}-order statistical tree if the covariance of any two of its leaf nodes depends only on their degree of shared evolution, i.e., on the tree-depth of their closest common ancestor:*

$$\mathrm{Cov}(V_{N,i}, V_{N,k}) =: c_n \tag{2}$$

where $n = \max\{b : \lfloor i/2^{N-b} \rfloor = \lfloor k/2^{N-b} \rfloor\}$, and $c_n, n = 0, \ldots, N$, is a fixed sequence of numbers.

Two binary tree models for network traffic, the wavelet independent Gaussian (WIG) [3] and multifractal wavelet model (MWM) [4] are 2^{nd}-order statistical trees. Both the WIG and MWM parsimoniously capture traffic properties like self-similarity and LRD and have fast $O(M)$ synthesis algorithms for M-point data sets. Note that while these models are 1^{st}-order stationary they are not 2^{nd}-order stationary due to their rigid tree structure.

Before solving the optimal sampling problem we define two terms.

Definition 2 *A tree process has a* positive correlation progression *if $c_m \geq c_{m-1}$ for $m = 1, \ldots, N-1$.*

Definition 3 *A tree process has a* negative correlation progression *if $c_m \leq c_{m-1}$ for $m = 1, \ldots, N-1$.*

3. OPTIMAL PROBING SCHEMES

In this section we consider the problem of determining which set of 2^p ($p \leq N$) leaf nodes gives the best LMMSE of the tree root R. In the traffic estimation scenario this corresponds to using 2^p probe packet-pairs with inter-spacing $\Delta = T/2^N$ to estimate R, with the constraint that the probes enter the queue at time instants equal to multiples of Δ.

Let $L = [l_i], i = 1, \ldots, 2^p$ be a vector of 2^p arbitrary but fixed leaf nodes at depth N. Denote by $S_L = [s_{i,j}]$ the covariance matrix of L; that is $s_{i,j} := \mathrm{Cov}(l_i, l_j)$.

Because of our assumption of a 2^{nd}-order statistical tree, the correlation between any leaf node and the root node is identical:

$$\mathrm{Cov}(l_i, R) = \rho. \tag{3}$$

It is well known [5] that the LMMSE of R given L (assuming zero-mean random variables) is given by

$$\widehat{R} = \rho \mathbf{1}_{1\times 2^k} S_L^{-1} L \tag{4}$$

and[1] that the resulting mean square error is

$$\mathcal{E} = \mathrm{var}(R) - \mathrm{Cov}(L,R)^T S_L^{-1} \mathrm{Cov}(L,R). \tag{5}$$

Due to the special form of the covariance vector $\mathrm{Cov}(L,R) = \rho \mathbf{1}_{1\times 2^k}$ we observe that minimizing $\mathcal{E}$ over all possible choices of sample leaves L is equivalent to maximizing the sum of all elements of the covariance matrix S_L^{-1}.

We define two special choices of L. The first set, called *bunch* samples, consists of all leaves belonging to a sub-tree with root at depth $N-p$ in the tree, that is,

$$\widehat{L} = \{V_{N,k} : \lfloor k/2^p \rfloor = K\}, \tag{6}$$

where the constant $K \in \{0, \ldots, 2^{N-p}-1\}$. The second set, called *uniform* samples, is such that each of its elements belongs to a different sub-tree rooted at depth p, that is,

$$L^* = \{V_{N,k}, k \in G\}, \tag{7}$$

where $G \subset \{0, \ldots, 2^N - 1\}$ such that $\forall i, j \in G$ and $i \neq j$, $\left\lfloor \frac{i}{2^{N-p}} \right\rfloor \neq \left\lfloor \frac{j}{2^{N-p}} \right\rfloor$.

In the case of positive correlation progression, leaf nodes closer to each other are more strongly correlated. Thus intuitively $\widehat{L}$ possesses the most redundant information about R among all choices L. We thus expect it to give the worst LMMSE. Conversely L^* should give the optimal LMMSE. A similar reasoning holds for the negatively correlated progression case with the rôles of $\widehat{L}$ and L^* reversed.

The following result vindicates our intuition regarding the worst-case sampling choices.

Theorem 4 *Assume that $\mathrm{var}(l_i | l_j, j \neq i) > 0, \forall i$, implying that S_L is positive definite $\forall L$. For a positively correlated progression, the sum of elements of S_L^{-1} is minimized by $L = \widehat{L}$. For a negatively correlated progression, the sum of elements of S_L^{-1} is minimized by $L = L^*$.*

The theorem is proved in the Appendix. We conjecture that a result similar to Theorem 4 holds for the case of the best-case sampling choice.

Conjecture 5 *For a tree with positive correlation progression, the sum of elements of S_L^{-1} is maximized by $L = L^*$. For a tree with negative correlation progression, the sum of elements of S_L^{-1} is maximized by $L = \widehat{L}$.*

This conjecture is supported by the experiments in Section 4.

4. EMPIRICAL COMPARISON OF PROBING SCHEMES

This section provides numerical support for Conjecture 5. We verify our claim using a WIG model of tree-depth 6 possessing an fractional Gaussian noise-like correlation structure corresponding to Hurst parameter $H = 0.8$ and $H = 0.3$. To be precise, we choose the WIG model parameters such that $\mathrm{var}(V_{j,k}) = 2^{-2jH} \cdot \text{constant}$ (see [3] for further details).

[1] Here $\mathbf{1}$ refers to a matrix with all elements equal to 1.

Note that $H > 0.5$ corresponds to positive correlation progression while $H \leq 0.5$ corresponds to negative correlation progression. With this correlation structure the WIG satisfies the assumption made in Theorem 4.

Fig. 4 compares the LMMSE error of the estimated root node $\widehat{R}$ (normalized by the variance of the root) of the uniform and bunch sampling patterns. Since an exhaustive search of all possible patterns is computationally expensive (for example there are over 10^{18} ways of choosing 32 leaf nodes from among 64) we instead compute the MSE for 10^4 randomly selected patterns. Observe that the bunch pattern gives the smallest LMMSE error for the negatively correlated process in Fig. 4(a) while the uniform pattern does so for the positively correlated one in Fig. 4(b), supporting our conjecture. As proved in Theorem 4, the bunch and uniform patterns give the worst LMMSE for the positive and negatively correlated processes respectively.

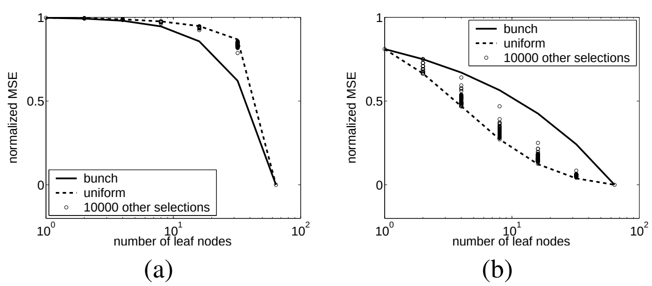

Fig. 4. *Comparison of probing schemes for a WIG model with (a) negative correlation progression and (b) positive correlation progression. Observe that the bunch is optimal in (a) while the uniform is optimal in (b). The uniform and the bunch give the worst performance in (a) and (b) respectively, which confirms our theoretical results.*

5. CONCLUSIONS

Our results have important implications for traffic estimation. The main finding of this paper is that the optimal probing scheme depends on the correlation structure of the traffic process. Internet traffic measurements typically reveal a positive correlation progression. Thus schemes like *cprobe* [6] that send a stream of back-to-back probe packets will give poor estimates of the average traffic rate. However in exceptional cases where the traffic could possess negative correlation progression, we conjecture that the same scheme will be optimal!

This paper solves only the "tip of the iceberg" of the larger problem of optimal sampling schemes on trees. What is the optimal set of M nodes, not necessarily leaf nodes, to estimate a specified node on a tree? How does the solution change if the tree is no longer a 2^{nd}-order statistical tree or even a binary tree? These are questions waiting to be explored and could have importance in several fields apart from networking.

Appendix

We now prove a Lemma which we then use to prove Theorem 4. As a first step we compute the leaf arrangements L which maximize and minimize the sum of all elements of S_L.

Lemma 6 *Assume a positive correlation progression. Then, $\sum_{i,j} s_{i,j}$ is minimized over all choices of L at $L = L^*$ and maximized at $L = \widehat{L}$. For a negative correlation progression, $\sum_{i,j} s_{i,j}$ is maximized over all choices of L at $L = L^*$ and minimized at $L = \widehat{L}$.*

Proof Let $h_m = \#\{s_{ij} \in S_L : s_{ij} = c_m\}$ be the number of elements of S_L equal to c_m. Define $a_m := \sum_{\alpha=0}^{m} h_\alpha, m \geq 0$ and set $a_{-1} = 0$. Then

$$\begin{aligned}
\sum_{i,j} s_{i,j} &= \sum_{m=0}^{N} c_m h_m = \sum_{m=0}^{N-1} c_m(a_m - a_{m-1}) + c_N h_N \\
&= \sum_{m=0}^{N-1} c_m a_m - \sum_{m=-1}^{N-2} c_{m+1} a_m + c_N h_N \\
&= \sum_{m=0}^{N-2} (c_m - c_{m+1}) a_m + c_{N-1} a_{N-1} - c_0 a_{-1} \\
&\qquad + c_N h_N \\
&= \sum_{m=0}^{N-2} (c_m - c_{m+1}) a_m + \text{constant}, \qquad (8)
\end{aligned}$$

where we used the fact that $a_{N-1} = a_N - h_N$ is a constant independent of the choice of L, since $h_N = 2^p$ and $a_N = 2^{2p}$.

We now show that $L = L^*$ maximizes $a_m, \forall m$ while $L = \widehat{L}$ minimizes $a_m, \forall m$. First we prove the results for $L = L^*$. Case (i) $m \geq p$. Since every element of L^* has distance at most $p - 1$ with all other elements, a_m is maximized by L^*. Case (ii) $m < p$ (assuming $p > 0$). Let the number of elements of L belonging to the sub-tree of $V_{m+1,q}$ be $g_q, q = 0, \ldots, 2^{m+1} - 1$. We have

$$a_m = \sum_{q=0}^{2^{m+1}-1} g_q(2^p - g_q) = 2^{2p-1+m} - \sum_{q=0}^{2^{m+1}-1} (g_q - 2^{p-1}) \qquad (9)$$

since every element of L in the sub-tree of $V_{m,q}$ must have distance *at most* m with all nodes *not* in the same sub-tree but must have distance *at least* $m + 1$ with all nodes *within* the same sub-tree.

The choice of g_q's is constrained to lie on the hyperplane $\sum_q g_q = 2^p$. Obviously the quadratic form of (9) is maximized by the point on this hyperplane closest to the point $(2^{p-1}, \ldots, 2^{p-1})$ which is $(2^{p-m-1}, \ldots, 2^{p-m-1})$. This is clearly achieved by L^*.

Now we prove the results for $L = \widehat{L}$.

Case (i) $m < N - p$. For $\widehat{L}$ we have $a_m = 0$, the smallest value it can take.

Case (ii) $N - p \leq m < N$. Consider node l_i which without any loss of generality belongs to the sub-tree of $V_{m+1,0}$. Let $a_m(l_i)$ be the number of elements of L to which l_i has distance less than or equal to m. Now since l_i has distance less than or equal to m only with those elements of L not in the same sub-tree, we must have $a_m(l_i) \geq 2^p - 2^{N-m-1}$. Since $\widehat{L}$

achieves this lower bound for $a_m(l_i), \forall i$ and $a_m = \sum_i a_m(l_i)$, $\widehat{L}$ minimizes a_m in turn. $\square$

Let us now study, to what extent the above results transfer to the actual matrix of interest S_L^{-1}. We start with a useful formula.

Lemma 7 *Denote the eigenvalues of S_L by λ_i. Assume that no leaf node of the tree can be expressed as a linear combination of other leaf nodes, implying that $\lambda_i > 0, \forall i$. Set $D_L = [d_{i,j}]_{2^p \times 2^p} := S_L^{-1}$. Then there exist positive numbers f_i with $f_1 + \ldots + f_p = 1$ such that*

$$\sum_{i,j=1}^{2^p} s_{i,j} = 2^p \sum_{i=1}^{2^p} f_i \lambda_i, \text{ and} \tag{10}$$

$$\sum_{i,j=1}^{2^p} d_{i,j} = 2^p \sum_{i=1}^{2^p} f_i / \lambda_i. \tag{11}$$

Furthermore, for both special cases, L^* and $\widehat{L}$, we may choose the weights f_i such that only one is non-zero.

Proof Since the matrix S_L is real and symmetric there exists an orthonormal eigenvector matrix $U = [u_{i,j}]$ that diagonalizes S_L, that is $S_L = U \Lambda U^T$ where Λ is diagonal with eigenvalues $\lambda_j, j = 1, \ldots, 2^p$. Define $w_j := \sum_i u_{i,j}$. Then

$$\begin{aligned}\sum_{i,j} s_{i,j} &= \mathbf{1}_{1\times 2^p} S_L \mathbf{1}_{2^p \times 1} \\ &= (\mathbf{1}_{1\times 2^p} U) \Lambda (\mathbf{1}_{1\times 2^p} U)^T \\ &= [w_1 \ldots w_{2^p}] \Lambda [w_1 \ldots w_{2^p}]^T \\ &= \sum_j \lambda_j w_j^2.\end{aligned}$$

Further, since $U^T = U^{-1}$ we have

$$\sum_j w_j^2 = (\mathbf{1}_{1\times 2^p} U)(U^T \mathbf{1}_{2^p\times 1}) = \mathbf{1}_{1\times 2^p} I \mathbf{1}_{2^p \times 1} = 2^p. \tag{12}$$

Setting $f_i = w_i^2/2^p$ establishes (10). Using the decomposition $S_L^{-1} = (U^T)^{-1}\Lambda^{-1}U^{-1} = U\Lambda^{-1}U^T$ similarly gives (11).

Consider the case $L = L^*$. Since L^* consists of a symmetrical set of leaf nodes (see (7)) the covariances of $l_i \in L^*$ with its fellow leaf nodes does not depend on i, and we can set:

$$\lambda^* := \sum_{j=1}^{2^p} s_{i,j} = c_N + \sum_{m=1}^{p} 2^{p-m} c_m. \tag{13}$$

With the sum of the elements of any row of S_{L^*} being identical, the vector $\mathbf{1}_{2^p\times 1}$ with equal coordinates is an eigenvector of S_{L^*} with eigenvalue λ^* equal to (13).

Recall that we can always choose a basis of orthogonal eigenvectors which includes $\mathbf{1}_{2^p\times 1}$ as the first basis vector. It is well known that the rows of the corresponding basis transformation matrix U will then be exactly these normalized eigenvectors. Since they are orthogonal to $\mathbf{1}_{2^p\times 1}$, the sum of their coordinates w_j ($j = 2, \ldots, 2^p$) must be zero. Thus, all f_i but f_1 vanish. (The last claim follows also from the observation that the sum of coordinates of the L^2-normalized $\mathbf{1}_{2^p\times 1}$ equals $w_1 = 2^p 2^{-p/2} = 2^{p/2}$; due to (12) $w_j = 0$ for all other j.)

The reasoning is similar for $S_{\widehat{L}}$, and we can define

$$\widehat{\lambda} := \sum_{j=1}^{2^p} s_{i,j} = c_N + \sum_{m=1}^{p} 2^m c_{N-m}. \tag{14}$$

$\square$

Proof of Theorem 4 Note that the weights f_i as well as the eigenvalues λ_i of Lemma 7 depend on the arrangement of the leaf nodes L. To avoid confusion, we denote by λ_i the eigenvalues of S_L for an arbitrary fixed set of leaf nodes L, and by λ^* and $\widehat{\lambda}$ the only relevant eigenvalues of L^* and $\widehat{L}$ according to (13) and (14).

Assume a positive correlation progression, and let L be an arbitrary set of leaf nodes. Lemma 6 and Lemma 7 then imply that

$$\lambda^* \;\le\; \sum_j \lambda_j f_j \;\le\; \widehat{\lambda}. \tag{15}$$

Since S_L is positive definite, we must have $\lambda_j > 0$. We may then interpret the middle expression as an expectation of the positive "random variable" λ with discrete law given by f_i. Jensen's inequality applies with the convex function $1/x$ ($x > 0$) and yields

$$\sum_j (1/\lambda_j) f_j \;\ge\; \frac{1}{\sum_j \lambda_j f_j} \;\ge\; \frac{1}{\widehat{\lambda}}. \tag{16}$$

In other words, $\sum_{i,j} d_{i,j}$ is minimized by $\widehat{L}$; that is, bunching the nodes at one end of the tree gives the worst error.

A similar argument holds for the case $c_m \le c_{m-1}$ which proves the Theorem. $\square$

6. REFERENCES

[1] W. Leland, M. Taqqu, W. Willinger, and D. Wilson, "On the self-similar nature of Ethernet traffic (extended version)," *IEEE/ACM Trans. Networking*, pp. 1–15, 1994.

[2] G. He and J. C. Hou, "On exploiting long-range dependency of network traffic in measuring cross-traffic on an end-to-end basis," *IEEE INFOCOM*, 2003.

[3] S. Ma and C. Ji, "Modeling video traffic in the wavelet domain," *Proc. of INFOCOM*, pp. 201–208, Mar. 1998.

[4] R. Riedi, M. S. Crouse, V. Ribeiro, and R. G. Baraniuk, "A multifractal wavelet model with application to TCP network traffic," *IEEE Trans. Info. Theory*, vol. 45, pp. 992–1018, April 1999.

[5] H. Stark and J. W. Woods, *Probability, random processes, and estimation theory for engineers*. Prentice-Hall, 1986.

[6] R. L. Carter and M. E. Crovella, "Measuring bottleneck link speed in packet-switched networks," *Performance Evaluation (Proceedings of Performance'96)*, vol. 27&28, pp. 297–318, 1996.