BAYESIAN NETWORK LOSS INFERENCE

Dong Guo and Xiaodong Wang

Department of Electrical Engineering
Columbia University
New York, NY 10027

ABSTRACT

In large-scale dynamic communication networks, end-systems can not rely on the network itself to cooperate in characterizing its own behavior. This has prompted research activities on methods for inferring internal network behavior based on the external end-to-end network measurements. In particular, knowledge of the link losses inside the network is important for network management. However it is impractical to directly measure packet losses or delays at every router. On the other hand, measuring end-to-end (from sources to receivers) losses is relatively easy. We formulate the problems of link in a network as Bayesian inference problems and develop several Markov chain Monte Carlo (MCMC) algorithms to solve them. We then apply the proposed link loss algorithms to data generated by the Network Simulator (NS2) software, and obtain good agreements between the theoretical results and the actual measurements.

1. INTRODUCTION

As the Internet grows in size and diversity, its internal behavior becomes more difficult to measure. On the other hand, many emerging applications, such as IP telephony, needs on-line information about the network internal links to support various QoS requirements. This has prompted some recent research activities on network tomography. In most of these works, the maximum likelihood (ML) estimator [1, 2] is employed as the inference tool. The expectation maximization (EM) algorithm is used in [3] to infer the network statistics with missing observations. However, when the network is large, the likelihood surface typically exhibits many local maxima which the ML estimator can easily be trapped into. Here we propose to approach the network inference problems using Monte Carlo Bayesian methods [4, 5], which are powerful global optimization techniques.

So far most existing works on loss inference focus on capturing the mean loss, whereas the inference of the loss distribution has not been addressed. However, it is noted in [6] that for the same mean loss rate, different loss patterns can produce different perceptions of QoS. Here we use a Gilbert model to characterize the bursty nature of the link packet loss process. We then develop Markov chain Monte Carlo (MCMC) algorithms to estimate both the mean loss rate and the transition probabilities in the Gilbert model, based on end-to-end packet loss measurements. Finally, we apply the link loss algorithms on data generated by the Network Simulator (NS2) software, and obtain good agreements between theoretical results and actual measurements.

The rest of this paper is organized as follows. In Section 2, we describe the link loss models employed in this paper. In Section 3, we present several MCMC algorithms for Bayesian inference of the link loss. In Section 4, we provide simulation results using data generated by NS 2. Section 5 contains the conclusions.

2. LINK LOSS MODELS

Link loss is usually caused by congestion (router buffer overflow), link failure and lossy links. Therefore, link loss measurements and models offer a better understanding of the network behavior.

Bernoulli model: The widely used link loss model is the simple Bernoulli model. That is, the link loss is modelled by a random process consisting of Bernoulli trials, in which the outcome of each experiment (e.g., a packet is either lost or delivered) is independent of previous trials. However, in practice, a packet loss is usually an indication of possible congestion buildup; and with high probability the next packet may also be lost, leading to a temporal dependency of link losses [7]. Such a dependency can be characterized by the Gilbert model [6].

Gilbert model: The Gilbert model is a two-state Markov model with parameters (p, q), where p is the probability that the next packet is lost, given the current packet is delivered; and q is the probability that the next packet is delivered, given the current one is lost. Note that typically we have $p + q < 1$. If $p + q = 1$, then the Gilbert model reduces to a Bernoulli model.

3. BAYESIAN INFERENCE OF LINK LOSS

In this section, we develop a number of algorithms for estimating the link loss parameters of a network based on end-to-end unicast or multicast loss measurements. We first consider estimating the mean link loss probabilities in a network under the Bernoulli model. Then we treat estimating the loss parameters in a network under the Gilbert burst loss model.

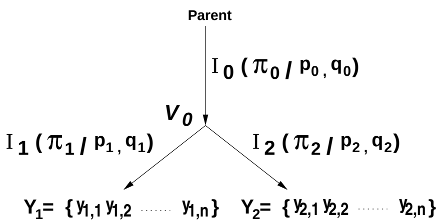

Fig. 1. The simple example network for mean loss ($\pi_i, i = 0, 1, 2$) inference or burst lost ($p_i, q_i, i = 0, 1, 2$) inference.

3.1. Mean Loss Inference

For the sake of clarity of presentation, we first illustrate our approach by using a simple example network shown in Fig. 1. Suppose the loss rates for the three links are π_0, π_1 and π_2, respectively. Our aim is to make a Bayesian inference about $\boldsymbol{\theta} \triangleq \{\pi_0, \pi_1, \pi_2\}$ based on end-to-end packet loss measurements $\boldsymbol{Y} = \{\boldsymbol{Y}_1, \boldsymbol{Y}_2\}$. Our approach is based on the Markov chain Monte Carlo method, and in particular, the Gibbs sampler. To that end, we define the following indicator variable I_i as the unknown observations at the internal mode V_0 in Fig. 1. Denote $\boldsymbol{I} \triangleq \{I_i,\ i = 1, \ldots, n\}$.

In order to implement a Gibbs sampler for this problem, we assign Beta priors to the link loss rates, i.e.,

$$\pi_k \sim \text{Beta}(a_k, b_k), \quad k = 0, 1, 2. \tag{1}$$

Then we have

$$p(\boldsymbol{\theta}, \boldsymbol{I}, \boldsymbol{Y}) \propto \prod_{i=1}^{n} p(y_{1,i} \mid \boldsymbol{\theta}, \boldsymbol{I}) p(y_{2,i} \mid \boldsymbol{\theta}, \boldsymbol{I}) p(\boldsymbol{\theta}, \boldsymbol{I}) \tag{2}$$

It then follows from (2) that the conditional posterior distribution of link loss rates are still Beta, given by

$$p(\pi_0 \mid \boldsymbol{Y}, \boldsymbol{I}, \pi_1, \pi_2) \sim \text{Beta}\left(a_0 + n - \sum I_i,\ b_0 + \sum I_i\right), \tag{3}$$

$$p(\pi_1 \mid \boldsymbol{Y}, \boldsymbol{I}, \pi_0, \pi_2) \sim \text{Beta}\left(a_1 + \sum I_i(1 - y_{1,i}),\ b_1 + \sum I_i y_{1,i}\right), \tag{4}$$

$$p(\pi_2 \mid \boldsymbol{Y}, \boldsymbol{I}, \pi_0, \pi_1) \sim \text{Beta}\left(a_2 + \sum I_i(1 - y_{2,i}),\ b_2 + \sum I_i y_{2,i}\right). \tag{5}$$

Moreover, the conditional posterior of the loss indicator I_i of the internal node is given by

$$p(I_i = 0 \mid \boldsymbol{Y}, \pi_0, \pi_1, \pi_2) \propto (1 - y_{1,i})(1 - y_{2,i})\pi_0, \tag{6}$$

$$p(I_i = 1 \mid \boldsymbol{Y}, \pi_0, \pi_1, \pi_2) \propto \pi_1^{1-y_{1,i}}(1 - \pi_1)^{y_{1,i}} \pi_2^{1-y_{2,i}} \times (1 - \pi_1)^{y_{2,i}}(1 - \pi_0). \tag{7}$$

The Gibbs sampler then iteratively draw random samples of $\{\pi_0, \pi_1, \pi_2, \boldsymbol{I}\}$ from the conditional marginal posterior densities (3)-(7).

For a general network, we can partition the set of nodes $\mathcal{V}$ as $\mathcal{V} = \mathcal{S} \bigcup \mathcal{D} \bigcup \mathcal{I}$, where $\mathcal{S}$ is the set of source nodes, $\mathcal{D}$ is the set of destination nodes, and $\mathcal{I}$ is the set of internal nodes. In order to estimate the mean link losses in the network based on end-to-end measurements, we define an indicator $\boldsymbol{I}_k = \{I_{k,i}, i = 1, \ldots, n\}$ for each node $k \in \mathcal{V}$, where $I_{k,i} = 1$ if packet i reaches node k, and $I_{k,i} = 0$ otherwise. Note that if $k \in \mathcal{S}$, then $I_{k,i} = 1$ for all i; and if $k \in \mathcal{D}$, then $I_{k,i} = y_{k,i}$ for each i.

For each node $k \notin \mathcal{S}$, define its parent $g(k)$ as the node that directly forwards packets to k. For each node $k \notin \mathcal{D}$, define its children $\mathcal{C}(k)$ as the set of nodes to which node k directly forwards packets. Denote the loss rate of the link between nodes ℓ and k as $\pi_{\ell,k}$. As before, denote the parameters of interest $\boldsymbol{\theta} = \{\pi_{g(k),k}, k \in \mathcal{V}\}$; the indicators $\boldsymbol{I} = \{\boldsymbol{I}_k, k \in \mathcal{I}\}$, where $\boldsymbol{I}_k = \{I_{k,i}, i = 1, \ldots, n\}$; and the packet loss measurements $\boldsymbol{Y} = \{\boldsymbol{Y}_k, k \in \mathcal{D}\}$, where $\boldsymbol{Y}_k = \{y_{k,i}, i = 1, \ldots, n\}$. Then we have the following algorithm for sampling from the joint posterior density $p(\boldsymbol{\theta}, \boldsymbol{I} | \boldsymbol{Y})$.

Algorithm 1 [Mean link loss inference]

Initialization: Draw random samples $\boldsymbol{\theta}^{(0)}$ and $\boldsymbol{I}^{(0)}$ from their priors.

For $j = 1, 2, \ldots, J$ (J is total number of samples)

- *Given $\boldsymbol{I}^{(j-1)}$, for each $k \in \mathcal{D} \bigcup \mathcal{I}$, draw a sample*

$$\pi_{g(k),k}^{(j)} \sim p\left(\pi_{g(k),k} \mid \boldsymbol{I}_k^{(j-1)}, \boldsymbol{I}_{g(k)}^{(j-1)}\right); \tag{8}$$

- *Given $\boldsymbol{\theta}^{(j)}$, for each $k \in \mathcal{I}$ and $i = 1, \ldots, n$, draw a sample*

$$I_{k,i}^{(j)} \sim p\left(I_{k,i} \mid \pi_{g(k),k}^{(j)}, I_{g(k),i}^{(j-1)}, \{\pi_{k,\ell}^{(j)}, I_{\ell,i}^{(j)}, \ell \in \mathcal{C}(k)\}\right). \tag{9}$$

Note that the conditional posterior densities (8)-(9) can be similarly calculated as in the simple example described above. Moreover, a sensible sampling strategy is a bottom-up approach where we start from the parent nodes of the destination nodes, followed by their grandparents nodes, and so on, until we reach the source nodes.

3.2. Bursty Loss Inference

As discussed in Section 2, the Gilbert model can capture the bursty nature of the link loss. Consider again the simple network depicted in Fig. 1. Now assume that each link loss is characterized by a Gilbert model with parameter (p_k, q_k), $k = 0, 1, 2$. Hence in this case the parameters of interest are $\boldsymbol{\theta} = \{p_0, q_0, p_1, q_1, p_2, q_2\}$. The observations $\boldsymbol{Y} = \{\boldsymbol{Y}_1, \boldsymbol{Y}_2\}$ as well and the indicators for the internal node $\boldsymbol{I}$ are defined the same as above. We next discuss methods for Bayesian inference of $\boldsymbol{\theta}$ based on $\boldsymbol{Y}$.

Denote $\boldsymbol{I}_{[-i]} = \{I_1, \ldots, I_{i-1}, I_{i+1}, \ldots, I_n\}$. Then the conditional posterior marginal distribution of I_i, $i = 1, \ldots, n$, is determined by

$$p(I_i = 0 \mid \boldsymbol{Y}, \boldsymbol{I}_{[-i]}, \boldsymbol{\theta}) = y_{1,i}\, y_{2,i} \Big[I_{i-1} p_1 + (1 - I_{i-1})(1 - q_1) \Big] \Big[I_{i+1} q_1 + (1 - I_{i+1})(1 - q_1) \Big], \quad (10)$$

$$p(I_i = 1 \mid \boldsymbol{Y}, \boldsymbol{I}_{[-i]}, \boldsymbol{\theta}) = \left[(1 - y_{1,i}) \frac{p_1}{p_1 + q_1} + y_{1,i} \frac{q_1}{p_1 + q_1} \right] \left[(1 - y_{2,i}) \frac{p_2}{p_2 + q_2} + y_{1,i} \frac{q_2}{p_2 + q_2} \right] \Big[I_{i-1}(1 - p_1) + (1 - I_{i-1}) q_1 \Big] \Big[I_{i+1}(1 - p_1) + (1 - I_{i+1}) p_1 \Big]. \quad (11)$$

On the other hand, the marginal posterior distributions $p(p_k | \boldsymbol{Y}, \boldsymbol{I}, \boldsymbol{\theta} \backslash p_k)$ [resp. $p(q_k | \boldsymbol{Y}, \boldsymbol{I}, \boldsymbol{\theta} \backslash q_k)$], $k = 0, 1, 2$, do not admit closed-form expressions. Here we use a Metropolis-Hastings step with a random walk increment proposal distribution, e.g.,

$$r\left(p_k \mid p_k^{(j-1)}\right) \sim \mathcal{N}\left(p_k^{(j-1)}, \sigma^2\right). \quad (12)$$

That is, we draw a sample $p_k^{(j)}$ from the above proposal distribution, and accept it with probability

$$\alpha\left(p_k, p_k^{(j-1)}\right) = \min\left\{1, \frac{p\left(p_k \mid \boldsymbol{Y}, \boldsymbol{I}^{(j-1)}, \boldsymbol{\theta} \backslash p_k\right)}{r\left(p_k \mid p_k^{(j-1)}\right)} \times \frac{r\left(p_k^{(j-1)} \mid p_k\right)}{p\left(p_k^{(j-1)} \mid \boldsymbol{Y}, \boldsymbol{I}^{(j-1)}, \boldsymbol{\theta} \backslash p_k\right)} \right\}, \quad (13)$$

where by assuming uniform prior on p_k, we have

$$p\left(p_0 \mid \boldsymbol{Y}, \boldsymbol{I}^{(j-1)}, \boldsymbol{\theta} \backslash p_0\right) \propto p\left(I_0^{(j)}\right) p_0^{n_{1,0}} (1 - p_0)^{n_{1,1}} \quad (14)$$

$$p\left(p_k \mid \boldsymbol{Y}, \boldsymbol{I}^{(j-1)}, \boldsymbol{\theta} \backslash p_k\right) \propto p\left(\tilde{y}_{k,0}^{(j-1)}\right) p_k^{n_{1,0}^{(k)}} (1 - p_k)^{n_{1,1}^{(k)}}, \quad k = 1, 2, \quad (15)$$

where $n_{u,v}$ is the number of occurrences of the adjacent pair (u, v) in the sequence $\boldsymbol{I}^{(j-1)}$, $u, v \in \{0, 1\}$; $n_{u,v}^{(k)}$ is the number of occurrences of the adjacent pair (u, v) in the subsequence $\tilde{\boldsymbol{Y}}_k$ of $\boldsymbol{Y}_k$, defined as $\tilde{\boldsymbol{Y}}_k = \{Y_{k,i} : I_i^{(j-1)} = 1\}$, $k = 1, 2$. Similarly we have

$$p\left(q_0 \mid \boldsymbol{Y}, \boldsymbol{I}^{(j-1)}, \boldsymbol{\theta} \backslash q_0\right) \propto p\left(I_0^{(j)}\right) q_0^{n_{0,1}} (1 - q_0)^{n_{0,0}} \quad (16)$$

$$p\left(q_k \mid \boldsymbol{Y}, \boldsymbol{I}^{(j-1)}, \boldsymbol{\theta} \backslash q_k\right) \propto p\left(\tilde{y}_{k,0}^{(j-1)}\right) q_k^{n_{0,1}^{(k)}} (1 - q_k)^{n_{0,0}^{(k)}}, \quad k = 1, 2. \quad (17)$$

The marginal distributions on the initial measurement are given by

$$p(I_0 = 0) = \frac{p_0}{p_0 + q_0} \quad (18)$$

$$\text{and} \quad p(\tilde{y}_{k,0} = 0) = \frac{p_k}{p_k + q_k}, \quad k = 1, 2. \quad (19)$$

The above simple random-walk proposal distribution is frequently used due to its simplicity. However, a small step size σ in the proposal distribution (12) will result in exceedingly slow movement of the corresponding Markov chain, whereas a large σ will result in very slow acceptance rate. To improve the efficiency of the sampling algorithm, here use the orientational bias Monte Carlo (OBMC) method [5], which enables an MCMC sampler to make large step-size jumps without lowering the acceptance rate.

Finally we summarize the link loss parameter estimation algorithm for the Gilbert loss model based on end-to-end packet loss measurements. We use the same notation as in Section 3.1. The parameters of interest are $\boldsymbol{\theta} = \{p_{g(k),k}, q_{g(k),k}, k \in \mathcal{V}\}$.

Algorithm 2 [Burst link loss inference]

Initialization: Draw random samples $\boldsymbol{\theta}^{(0)}$ *and* $\boldsymbol{I}^{(0)}$ *from their priors.*

For $j = 1, 2, \ldots, J$ *(*J *is total number of samples)*

- *Given* $\boldsymbol{\theta}^{(j-1)}$, *for each* $k \in \mathcal{I}$ *and* $i = 1, \ldots, n$, *Draw a sample*

$$I_{k,i}^{(j)} \sim p\Big(I_{k,i} \mid p_{g(k),k}^{(j-1)}, q_{g(k),k}^{(j-1)}, \{I_{g(k),m}^{(j-1)}\}, \{p_{k,\ell}^{(j-1)}, q_{k,\ell}^{(j-1)}, I_{\ell,m}^{(j-1)}, \ell \in \mathcal{C}(k)\}, m = i-1, i, i+1\}\Big). \quad (20)$$

- *Given* $\boldsymbol{I}^{(j)}$, *sample* $\boldsymbol{\theta}^{(j)}$ *using the OBMC method as discussed above.*

4. SIMULATIONS

In this section, we provide simulation results on a NS2 platform [8] to illustrate the good performance of the proposed MCMC Bayesian network inference methods. As shown in Fig. 2, network with 8 nodes and 7 links is used in our simulation. The link loss probability π_k for the seven links are set to be $\{0.01, 0.001, 0.01, 0.02, 0.001, 0.02, 0.01\}$, respectively. In order to model the burst link loss, we use the Gilbert model to model the loss at the seven links, with parameters (p_k, q_k) as $\{(0.01, 0.85)$, $(0.001, 0.95)$, $(0.01, 0.85)$, $(0.01, 0.85)$, $(0.01, 0.85)$, $(0.001, 0.85)$,$(0.01, 0.95)$,$(0.01, 0.85)\}$, respectively.

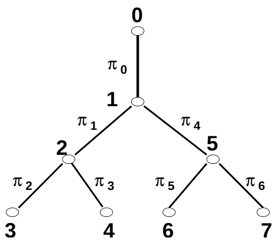

Fig. 2. The simulated network by NS2.

We consider the performance of the MCMC estimators for link loss using data collected from NS2. In general, for each simulation, $J_0 + J = 500$ samples are drawn in the MCMC procedure with the first $J_0 = 300$ samples as "burn-in" period and discarded. In the OBMC algorithm, the number of reference point is set to be $K = 5$.

The boxplot of the MCMC samples for the Bernoulli random loss model is shown in Fig. 3. In the mean loss inference, 1000 measurements from NS2 are used. It is seen that the mean of these samples is very close to the true values. The maximum error of the mean link loss estimation is only about 0.001. Moreover, these samples closely concentrate on the small region near the true values.

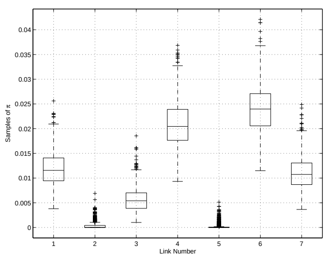

Fig. 3. The box plot for mean link loss inference under the Bernoulli model. .

Next we consider the inference of loss parameters under the Gilbert burst loss model. In this experiments, 5000 measurements from NS2 are used. The boxplot for the burst loss inference is shown in Fig. 4. In general, the algorithms give good estimates of the loss parameters. However, if the parameter has a very small value, e.g., $p = 0.001$ at link 2 and link 5, and with insufficient number of measurement data, these algorithms will be trapped to some local stationary points. The reason for this is that in this case it is very difficult to get enough samples to make inference on these parameters with extremely small values.

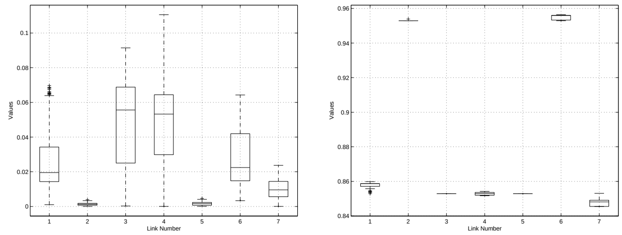

Fig. 4. The box plot for burst link loss inference under the Gilbert model.

5. CONCLUSION

In this paper, several Markov chain Monte Carlo (MCMC) algorithms are developed to solve the problems link loss estimation in a network based on end-to-end measurements. With a Bernoulli link loss model, a Gibbs sampler is given for inferring the link loss behavior. For the Gilbert bursty loss model, which can provide a more realistic characterization of the network loss, we propose a method for inferring the loss parameters based the Metropolis-Hasting method. The simulations show the good performance of the proposed MCMC algorithms.

6. REFERENCES

[1] R. Caceres, N. Duffield, J. Horowitz, and D. Towsley, "Multicast-based inference of network-internal loss characteristics," *IEEE Tran. Infom. Theory.*, vol. 45, no. 7, pp. 2462–2480, May. 2002.

[2] M. Coates and R. D. Nowak, "Sequential Monte Carlo inference of internal delays in nonstationary communication networks," *IEEE Trans. Sign. Proc.*, vol. 50, no. 2, pp. 366–376, Feb. 2002.

[3] C. Ji and A. Elwalid, "Measurement-based network monitoring and inference: scalability and missing information," *IEEE J. Select. Areas Commun.*, vol. 20, no. 4, pp. 714–725, May. 2002.

[4] W.R. Gilks, S. Richardson, and D.J. Spiegelhalter, *Markov chain Monte Carlo in Practice*, Chapman & Hall, 1995.

[5] J.S. Liu, *Monte Carlo Strategies for Scientific Computing*, Springer-Verlag, New York, 2001.

[6] W. Jiang and H. Schulzrinne, "QoS measurement of internet real-time mmultimedia services," in *Technical Report CUCS-015-99, Department of Computer Science, Columbia University*, 1999.

[7] V. Paxson, "End-to-end internet packet dynamics," *IEEE/ACM Trans. Networking*, vol. 7, no. 3, pp. 277–292, June 1999.

[8] "The network simulator-NS2," http://www.isi.edu/nsnam/ns/.