Adjustment of the BART Kalman Filter to Improve Real-Time Estimation of End-to-End Available Bandwidth

Erik Hartikainen[‡¶], Svante Ekelin[‡], and Johan M Karlsson[¶]

[‡] Ericsson Research
[¶] Department of Science and Technology, Linköping University

***Abstract*—This paper concerns simulation results using the network simulator ns-2 to validate the end-to-end available bandwidth estimation ability of the real-time measurement tool BART (Bandwidth Available in Real-Time). In addition to an active probing scheme, this method applies Kalman filtering in order to produce available bandwidth estimates. Simulation results show that clever use of an adjustable filter parameter yields great possibilities of achieving highly accurate estimates, even in the presence of harsh circumstances.**

***Keywords*—active probing, available bandwidth, Kalman filter, ns-2 simulation, real-time estimation**

I. INTRODUCTION

In recent years, there has been a rapid progress in the field of measuring available bandwidth. By using an active probing technique, and thereafter analyzing the effects that appear when cross traffic interferes with probe traffic, it is possible to estimate the available bandwidth on a network path.

In [1] Ekelin and Nilsson proposed the possibility of making use of Kalman filtering when analyzing the received probe traffic. The method, which was given the name BART (Bandwidth Available in Real-Time), applies an active probing scheme similar to the TOPP method [2]. By transmitting trains of probe packets at randomized data rates, it is possible to cause self-induced congestion along a network path. This characteristic makes it possible to utilize the measure inter-packet strain, i.e. the relative increase in the time separation of consecutive probe packets. The strain is a fundamental input parameter to the Kalman filter in BART.

The main novelty of BART is the technique for using Kalman filtering [3] in order to produce estimates of the available bandwidth in real-time. Considering already developed tools performing available bandwidth measurements, the majority require a considerable time for both probing and analysis before an estimate is presented. BART, on the other hand, produces estimates continuously as the trains of probe packets reach the receiver. In addition, the filtering approach has a number of promising features, making it possible to tune BART according to specific needs of a certain measurement application. Although only a few parameters need manual adjustment, it is, for example, feasible to trade accuracy for agility in case of sudden changes of the network properties.

In this paper, the purpose is to investigate the BART method in a simulation environment. More specifically, the network simulator ns-2 is used and the simulations stress the behavior of BART in scenarios with abrupt changes in cross traffic intensity and network link characteristics. The analysis of received probe traffic will highlight the effect of tuning a filter parameter with respect to existing circumstances. Experiments in a laboratory network and over the Internet have been performed in [4]; however, those experiments do not examine the potential of BART concerning adjustments of the Kalman filter parameters.

As a benchmark to the performance of BART, the tool pathChirp [5] has been used in the ns-2 environment. Among available measurement tools, e.g. [2, 5-8], pathChirp is probably the method that comes closest to offer similar functionality as BART; since the aim of pathChirp is to continuously estimate the end-to-end available bandwidth.

II. MEASURING AVAILABLE BANDWIDTH USING BART

A. *Definition of available bandwidth*

Throughout this paper, the definition of available bandwidth can be summarized as:

$$B = \min_{j} (C_j - X_j) \tag{1}$$

The above equation describes that each link j along a network path has a given capacity C_j, which is determined by the network interfaces in the nodes connected to the specific link j. Typically, the capacity of each link j is not changeable whereas the opposite holds concerning the link load. Each link j can be affected by cross traffic X_j that most likely varies with time.

With respect to the link capacity C_j and the present cross traffic intensity X_j, it is possible to determine the available bandwidth of link j as $B_j = C_j - X_j$. The link that provides the smallest value of available bandwidth along the network path becomes the bottleneck link, which also determines the minimum available bandwidth B of the entire path. This argument is in accordance with the bulk of previous work on bandwidth estimation; e.g., consider [2, 5-10].

B. *Inter-packet strain*

BART is designed with respect to a network model that considers a network path as a series of first-come-first-served (FCFS) hops, where each hop includes a first-in-first-out (FIFO) queue along with a transmission link. In presence of several active connections causing an aggregated traffic situation, this implies that all traffic streams receive a share of the link capacity that is proportional to their intensity on the particular link. Further, the estimates of BART are based on information regarding the inter-packet strain of received probe packets. The strategy is to momentarily cause an overload situation on the network path by transmitting probe packets at a given rate u. It was shown in [2] that variations of u enable identification of the threshold for congestion, which provides useful information when producing an available bandwidth estimate of the network path.

The relation between the probe traffic intensity and the available bandwidth can be described by the dimensionless quantity inter-packet strain. For a probe traffic intensity $u \leq B$, the expectation value of the strain equals zero as there is no congestion on the network path. However, in case $u > B$, congestion occurs and in the region where only one link is overloaded, the strain has been approximated using a fluid model in [2] to grow linearly with respect to the probe traffic overload. See Figure 1.

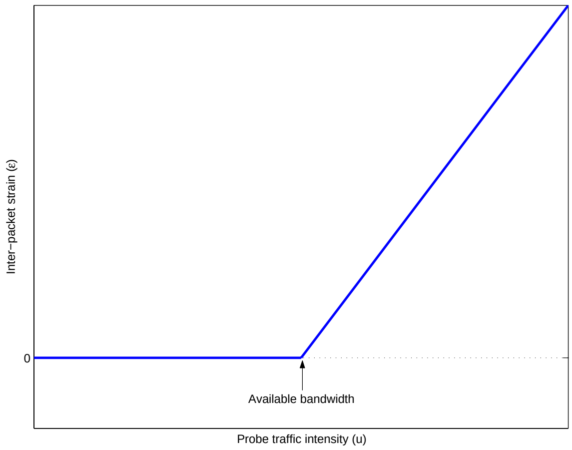

Figure 1. Information regarding the inter-packet strain and the probe traffic intensity makes it possible to estimate the available bandwidth.

By having knowledge of the time separation at the transmitter and by time-stamping the probe packets at the receiver, the average inter-packet strain ε may be computed as:

$$\varepsilon = \frac{1}{N}\sum_{i=2}^{N+1}\frac{(\tau_i^* - \tau_{i-1}^*) - (\tau_i - \tau_{i-1})}{(\tau_i - \tau_{i-1})} \tag{2}$$

The time probe packet i leaves the transmitter is denoted as τ_i whereas τ_i^* is the arrival time at the receiver. Hence, the relative inter-packet strain for N consecutive packet pairs is computed when the received probe train consists of $N + 1$ packets.

In addition to the average inter-packet strain ε, its variance is also computed for each received probe train:

$$R = \frac{1}{N(N-1)}\sum_{i=2}^{N+1}\left(\frac{(\tau_i^* - \tau_{i-1}^*) - (\tau_i - \tau_{i-1})}{(\tau_i - \tau_{i-1})} - \varepsilon\right)^2 \tag{3}$$

The variance R is of great importance for the available bandwidth estimation, since it serves as an estimator for the precision of the inter-packet strain measurement.

C. *Estimation using Kalman filtering*

In order to apply a Kalman filter formalism, the strain measurements are described by a piecewise linear model:

$$\varepsilon = v + \begin{cases} 0 & (u \leq B) \\ \alpha u + \beta & (u > B) \end{cases} \tag{4}$$

In the above expression, the theoretical model in Figure 1 is complemented by a variable v, the measurement noise. α and β are the state variables of the system, which describe the state vector $x = [\alpha\ \beta]^{\mathrm{T}}$. The state vector contains all the information needed to express the sloping straight line in the theoretical model, illustrating the inter-packet strain in case of congestion. Hence, since the Kalman filter is estimating the state vector, an estimate of the available bandwidth may be computed as (caret marks indicate estimates):

$$\hat{B} = -\frac{\hat{\beta}}{\hat{\alpha}} \tag{5}$$

Because of the applied strain model being only piecewise linear, an ordinary execution of the Kalman filter would not work. It needs to operate in a linear region. Thus, it is of great importance to only make use of reliable strain measurements affected by congestion. One approach is to neglect all measurements except those for which the probe traffic intensity is larger than the current estimate of the available bandwidth, suggested in [1, 4]. This will most likely regulate the Kalman filter to only be working in the region of congestion, which also is the area of interest for making available bandwidth estimates.

Two important parameters for the performance of the BART Kalman filter are the process noise covariance matrix and the variance of the already mentioned measurement noise. Whereas the measurement noise component represents deviations in strain measurements that do not reflect the true traffic situation, the process noise accounts for fluctuations in the system causing changes to the state vector.

In the application of BART, the variance of the measurement noise is computed as the variance R of the average inter-packet strain ε, see (3). The process noise covariance matrix Q, on the other hand, needs to be supplied. This makes it possible to tune BART for better performance, provided some information of the system's characteristics is known beforehand.

Q can be related to anything that affects the state of the system, e.g., statistical properties of the cross traffic and variations in the link capacity when switching bottleneck link along the network path. An alternative is to simply consider Q as an adjustable parameter, which determines the trade-off between quick adaptation and accuracy. The reason for this is that Kalman filtering continuously involves comparisons of the prediction error covariance P (which is directly affected by Q) and the measurement noise variance R. This is done in such a way as to find the optimal balance between the prediction from the previous estimate and the correction from the new measurement.

For a more comprehensive description of BART and the Kalman filter equations, the reader is referred to [1, 4] and [3] respectively.

III. SIMULATION MODEL

Simulations have been carried out using the network simulator ns-2. The focus has been on investigating the behavior of BART as abrupt changes occur in the system. For a schematic view of the used network topology, consider Figure 2.

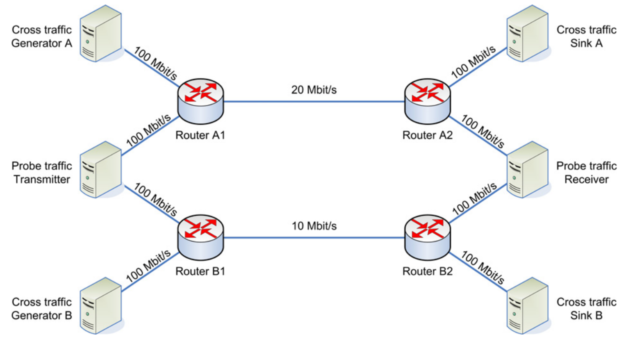

Figure 2. An illustration of the network topology used in the ns-2 environment.

In the simulations, the probe traffic has been provided with the possibility of using two different routes to reach its destination. By using path A, including router A1 and A2, the bottleneck link capacity is 20 Mbit/s. The alternative route is path B, which makes use of router B1 and B2, with a bottleneck link capacity of 10 Mbit/s.

Simulations have been performed with several cross traffic constellations. However, this paper only presents results of a scenario where the cross traffic generators simulate aggregated UDP traffic caused by 10 active users. Each user transmits packets at intervals following a Pareto distribution (shape parameter = 1.9), a tail-heavy distribution with infinite variance. Table I shows the distribution of the packet sizes, which roughly corresponds to observations from Sprint[1]. Further, the generators are configured such that each user approximately utilizes an equivalent transmission rate. Cross traffic leaving generator A is not affecting path B, and vice versa.

TABLE I
DISTRIBUTION OF CROSS TRAFFIC PACKET SIZES IN BYTES

Packet size:	Percentage:
40 B	30.23 %
52 B	18.75 %
576 B	10.10 %
1420 B	15.15 %
1500 B	25.77 %

BART is configured to use probe packets of 1500 bytes and probe trains including 17 packets per train. The probe traffic intensity of each train is randomly chosen from a uniform distribution, covering the interval from 1 Mbit/s to 20 Mbit/s. The inter-departure time between two trains is set to one second. This entails that BART injects probe traffic with an average intensity of 0.204 Mbit/s.

When utilizing pathChirp, the default parameters included in the available code have been used[2]. However, according to a comment in the code, initial values need not to be set, since pathChirp is expected to adapt to the present situation. An examination of the performed pathChirp simulation reveals that approximately 1.95 estimates are produced every second. The average probe traffic intensity was measured to be 0.300 Mbit/s.

In addition to the estimates of BART and pathChirp, the actual available bandwidth is measured as a reference. This is accomplished by subtracting the average cross traffic intensity from the bottleneck link capacity. The average is calculated using a sliding window of three seconds (this value is chosen to suit pathChirp).

The used simulation scenario covering two rapid changes in the system can be described as follows:

1. *Within the time interval $0 \leq t < 500$ s:*
   - Probe traffic is transmitted through path A, including router A1 and A2.
   - Generator A produces cross traffic with an average intensity of 15.0 Mbit/s.
   - The bottleneck link capacity is 20.0 Mbit/s.
2. *Within the time interval $500 \leq t < 1000$ s:*
   - Probe traffic is transmitted through path B, including router B1 and B2.
   - Generator B produces cross traffic with an average intensity of 7.5 Mbit/s.
   - The bottleneck link capacity is 10.0 Mbit/s.
3. *Within the time interval $1000 \leq t < 1500$ s:*
   - Probe traffic is transmitted through path B, including router B1 and B2.
   - Generator B produces cross traffic with an average intensity of 5.0 Mbit/s.
   - The bottleneck link capacity is 10.0 Mbit/s.

According to the theoretical sloping straight line model in Figure 1, the sudden change after 500 seconds corresponds to a step change in the state variable α, which is equal to the inverse of the bottleneck link capacity, while β is unchanged. Likewise, as the cross traffic intensity decreases after 1000 seconds, there is a step change in β, while α is unchanged.

IV. SIMULATION RESULTS

Previous evaluations of BART have not explored tuning of Q for improved performance. This section will highlight the behavior of the method using different Q.

In BART, Q is a symmetric 2×2-matrix consisting of three independent scalar elements. An appropriate illustration, showing the significance of applying different Q, can be seen in Figure 3.

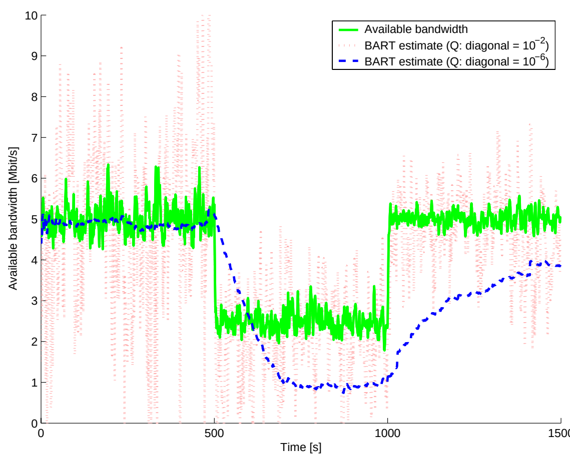

Figure 3. Available bandwidth estimates illustrating the trade-off between accuracy and agility in scenarios with sudden changes in the network.

[1] http://ipmon.sprintlabs.com/packstat/packetoverview.php (21 Sept. 2005)
[2] http://www.spin.rice.edu/Software/pathChirp/ (21 Sept. 2005)

Figure 3 depicts results using two rather extreme cases of the Q matrix. The dashed curve shows the performance of BART when the Q matrix is taking the relatively small value of 10^{-6} on the main diagonal and zeroes for the other elements. This choice of Q is probably fairly convenient for networks with slowly varying characteristics; however, the adaptation is very slow in presence of abrupt variations. The dotted curve illustrates the opposite feature; in this case the diagonal of the Q matrix is assigned 10^{-2}, which is a rather high value. The agility of the filter is the main advantage with this choice of Q, but the estimates contain a lot of measurement noise, causing a high variance.

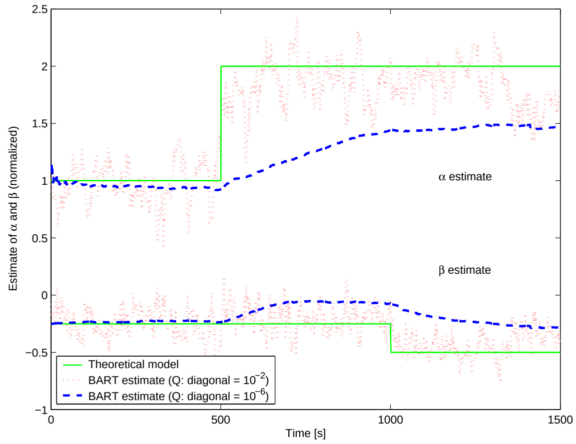

Figure 4. Estimating the state variables using process noise covariance matrices with different characteristics.

Figure 4 demonstrates how BART is estimating the state variables using the two different Q matrices. Again, it is clear that one approach has a slowly varying behavior, rather insensitive to the measurement noise, whereas the opposite holds for the other configuration. For the curve representing fast adaptation, BART is quick to realize when to change its estimate of a certain state variable; α when switching bottleneck link capacity after 500 seconds and β after 1000 seconds due to a different cross traffic intensity. The slower configuration needs a considerable time before adjusting to the new circumstances. Observe that the theoretical model refers to values corresponding to the sloping straight line, see Figure 1. Thus, since it is just an estimated model, it does not necessarily coincide with the simulation results. The values are normalized with respect to the maximum probe traffic intensity.

As a comparison to the results above, it is of interest to investigate the potential of BART when allowing “on the fly” adjustments of Q. An identical simulation as previously, in terms of cross traffic intensities and bottleneck link capacities, has been performed using a Q matrix that adjusts its values according to situations known beforehand.

When the system is expected to be in a steady state, BART is configured to put plenty of trust in its own ability of predicting the available bandwidth. This is accomplished by describing the process noise covariance with low values, e.g.:

$$Q = \begin{bmatrix} 10^{-6} & 0 \\ 0 & 10^{-6} \end{bmatrix} \qquad (6)$$

When sudden changes are likely to occur in the system, BART should not rely as much on its predicting ability. Instead, it should give greater weight to recent measurements by making appropriate adjustments of the elements in Q.

In case the new situation only requires tuning of α, an appropriate description of the process noise covariance could be:

$$Q = \begin{bmatrix} 10^{-2} & 0 \\ 0 & 10^{-6} \end{bmatrix} \qquad (7)$$

On the other hand, if only β is expected to change, a fairly decent choice of Q could be:

$$Q = \begin{bmatrix} 10^{-6} & 0 \\ 0 & 10^{-2} \end{bmatrix} \qquad (8)$$

By intelligent use of the three Q matrices suggested above, BART can produce fast and accurate estimates; consider Figure 5.

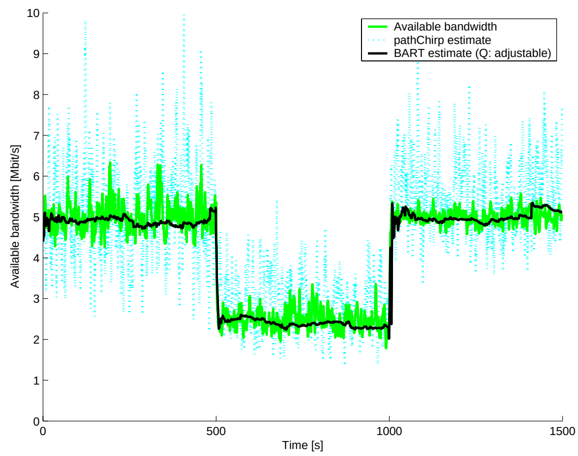

Figure 5. Available bandwidth estimates produced by BART (illustrating the potential of the method to leverage on additional information by adapting the Q matrix) and pathChrip.

In general, information of rapid changes in dynamic systems is not known in advance. Hence, the above adjustment approach of Q is not possible in all circumstances. Nevertheless, it is obvious from Figure 5 that BART has a great potential in producing high quality estimates of the available bandwidth, as long as the Kalman filter can be provided with suitable information.

As a benchmark to the performance of BART, Figure 5 also depicts the behavior of pathChirp. Although pathChirp shows an impressive agility, the variance of the estimates is very high. The latter remark is in line with observations from measurements made in [4].

Concerning the estimates of the state variables when using online adjustments of Q, see Figure 6.

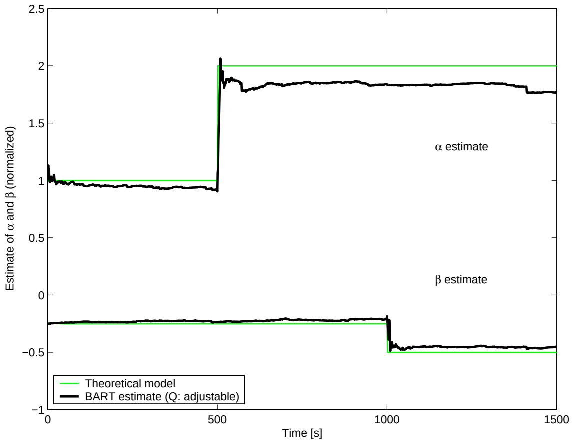

Figure 6. Estimating the state variables using an online adjustment approach of the process noise covariance matrix.

According to Figure 6, the performance of BART is very promising; fast reactions in case of a change in the system and small variations giving the impression of certainty. Although BART provides accurate estimates in Figure 5, a deviation from the theoretical model can be seen in Figure 6. It could be due to an imperfect theoretical model; this is currently under investigation.

V. Conclusion

The objective of this work was to investigate the performance of BART when abrupt changes occur in a communication network. Simulations performed in the ns-2 environment showed that accuracy and agility of the available bandwidth estimates are greatly dependent on the process noise covariance.

A tuning approach was demonstrated, yielding both quick adaptation and accurate estimates as BART was provided with a suitable filter parameter regulation. BART has been shown to possess a great potential for producing high quality estimates of the available bandwidth. The future work will involve various extensions of the method.

References

[1] S. Ekelin and M. Nilsson, "Continuous monitoring of available bandwidth over a network path," in *Proc. 2nd Swedish National Computer Networking Workshop*, Karlstad, Sweden, 2004.

[2] B. Melander, M. Björkman, and P. Gunningberg, "A new end-to-end probing and analysis method for estimating bandwidth bottlenecks," in *Proc. IEEE Globecomm '00*, San Francisco, USA, 2000.

[3] G. Bishop and G. Welch, "An introduction to the Kalman filter," in *SIGGRAPH 2001 Course 8*, 2001.

[4] S. Ekelin, M. Nilsson, E. Hartikainen, A. Johnsson, J.-E. Mångs, B. Melander, and M. Björkman, "Real-time measurement of end-to-end available bandwidth using Kalman filtering," Under submission.

[5] V. Ribeiro, R. Riedi, R. Baraniuk, J. Navratil, and L. Cottrell, "pathChirp: efficient available bandwidth estimation for network paths," in *Proc. Passive and Active Measurement Workshop*, 2003.

[6] M. Jain and C. Dovrolis, "Pathload: a measurement tool for end-to-end available bandwidth," in *Proc. Passive and Active Measurement Workshop*, 2002.

[7] J. Strauss, D. Katabi, and F. Kaashoek, "A measurement study of available bandwidth estimation tools," in *Proc. ACM SIGCOMM Internet Measurement Conference*, 2003.

[8] N. Hu and P. Steenkiste, "Evaluation and characterization of available bandwidth probing techniques," in *IEEE JSAC Internet and WWW measurement, mapping, and modeling*, 2003.

[9] M. Jain and C. Dovrolis, "End-to-end available bandwidth: measurement methodology, dynamics, and relation with TCP throughput," in *Proc. ACM SIGCOMM*, 2002.

[10] R. Prasad, M. Murray, C. Dovrolis, and K. Claffy, "Bandwidth estimation: metrics, measurement techniques, and tools," in *IEEE Network Magazine*, 2003.