

Optimal on-line estimation of the size of a dynamic multicast group

Sara Alouf, Eitan Altman and Philippe Nain

INRIA Sophia Antipolis, B.P. 93
06902, Sophia Antipolis Cedex, France
{salouf, altman, nain}@sophia.inria.fr

Abstract—In this paper we propose an efficient on-line estimation algorithm for determining the size of a dynamic multicast group. By using diffusion approximation and Kalman filter, we derive an estimator that minimizes the mean square of the estimation error. As opposed to previous studies, where the size of the multicast group is supposed to be fixed throughout the estimation procedure, we consider a dynamic estimation scheme that updates the estimation at every observation step. The robustness of our estimator to violation of the assumptions under which it has been derived is addressed via simulations. Further validations of our approach are carried out on real audio traces.

I. INTRODUCTION

SEVERAL papers have recently addressed the important question of estimating a population size in a multicast tree: a source might be interested to know how many recipients are connected to the multicast session (or are actively following some application that is being broadcast). When such an information is required, the source could ask all connected members to send an acknowledgment (ACK). But this could be undesirable in case of large populations as the acknowledgments could overload the network. To avoid this, one could alternatively ask each active connected receiver to send an acknowledgment with some small probability p^* . Yet, if p^* is chosen too small, the estimation would be inaccurate.

If we wish to further have the possibility of tracking the population size, we should repeat occasionally the requests for ACKs, at say, every S seconds. If S is not too large then the population size at two consecutive estimation instants would present statistical dependence. The engineering question we pose in this paper is how can we benefit from this dependence in order to be able to get better estimation, or alternatively, to get a given quality of estimation with a smaller required volume of ACKs (i.e. decreasing p^* or increasing S). A precise mathematical formulation of this problem would require to use the theory of nonlinear stochastic filtering, which does not provide us with tractable solutions. We shall instead use some simplifying assumptions that will allow us to obtain a good estimation scheme which, even if it is not always the optimal, will show good performance. To that end we shall make some simplifying assumptions: we shall consider an exponential distribution for

the time during which a receiver stays in the multicast session and make a large group size assumption. This will allow us to obtain a diffusion approximation for the dynamics of the population. Sampling this process at some regular time intervals will yield a discrete-time linear stochastic difference equation for the population dynamics. We will further derive a linear discrete-time equation for the measurements. The fact that both the population dynamics and the measurements in our approximations are linear, will allow us to use the powerful Kalman filtering theory to design simple dynamic estimation procedures which are optimal for the heavy traffic model (in minimizing the second moment of the error). These schemes thus make the best use of previous estimation in order to update the current estimation optimally.

Having proposed a dynamic estimation procedure that is optimal in our simplified mathematical model, we will test it on real traces which do not satisfy the assumptions of that model. Nevertheless, we report good performance of this procedure in these cases.

Estimating the size of a multicast session can be quite useful to many protocols like RTP [1], [2] and SRM [3] requiring such estimates for feedback suppression, as well as to several applications. For instance, Bolot, Turletti and Wakeman [4] use a probabilistic scheme in order to estimate the number of receivers in a multicast group and further estimate the proportion of congested receivers. Their mechanism is implemented in the IVS¹ videoconference system. In another context, Nonnenmacher and Biersack [5], [6] investigate the scalability of feedbacks as needed for reliable multicast and for the estimation of the number of receivers. They derive a random delay response scheme that scales well to group sizes as large as 10^6 receivers. The feedback implosion problem is handled at the receivers: each participant multicasts its response unless he receives one from another participant. In [7], Friedman and Towsley address the issue of estimating multicast session membership size. Mapping the polling mechanism to the problem of estimating the parameter n of the binomial (n, p) distribution, they derive an interval estimator for n and bounds for the amount of feedback as well as the polling probability in order to achieve specific requirements. They apply their results on both mechanisms introduced in [4], [5], [6] which have point

Part of this work was done when E. Altman was at C.E.S.I.M.O. Centro de Simulacion y Modelos Universidad de Los Andes (U.L.A.), Facultad de Ingenieria, Mérida, Venezuela.

¹available at <http://www-sop.inria.fr/rodeo/ivs.html>

estimators and further add some contributions to each. Another timer-based feedback scheme is proposed in [8] where receivers send their randomly delayed response only to the source which in turn initiates a new round of replies. Each request for replies sent by the source would reset the timers at the receivers. Two versions of the mechanism are proposed depending on whether the estimation is based on the first arrival solely or on all the received responses. The latter version improves the accuracy of the estimator but in both cases, the probability of a feedback implosion is not negligible.

A. Motivation

A naive approach to the estimation problem would consist in dividing the total number of ACKs received at the n th observation step by the ACK probability p^* . This ratio is then an estimator of the size of the multicast group at time nS . It is expected to have this estimator perform very poorly both because it does not take into account the “history” of the membership process and because the ratio of the number of ACKs received over the group size will converge to p^* a.s. only when the group size is large (the strong law of large numbers).

Our experiments reported in Fig. 1 have confirmed the poor behavior of this naive estimator. Fig. 1 depicts the evolution of membership size of an audio session recorded in December 1996 (details in Section V) and the estimation returned by the naive approach. The ACK interval S is set to 1 second. Two values (0.01 and 0.5) for the ACK probability p^* were retained. Fig. 1(a) displays the evolution of the number of group members and its “naive” estimation over time for $p^* = 0.01$. Obviously, there is too much noise in the estimation. Results for $p^* = 0.5$ are shown in Fig. 1(b). Even for this high (undesirable) probability, the naive estimator performs poorly. It is clear that some filtering is necessary in order to reduce the estimation error which motivates this work.

The rest of the paper is organized as follows: the mathematical model is introduced in Section II and the optimization is performed in Section III. The robustness of the estimator is addressed both via simulations (in Section IV) and via validations with real audio traces (in Section V). Open issues regarding the “optimal” selection of p^* and S are discussed in Section VI and concluding remarks are given in Section VII.

II. THE MODEL

We consider a multicast group that participants join and leave at random times. Let T_i and $T_i + \sigma_i$ be the join time and the leave time, respectively, of the i -th participant. In the following, $\sigma_i > 0$ is called the on-time of the i -th participant and $\{\sigma_i\}_i$ is referred to as the on-time sequence. Let $N(t)$ be the number of participants at time t or, equivalently, the size of the multicast session at time t . Under the enforced assumption that $N(0) = 0$, we have

$$N(t) = \sum_i \mathbf{1}(T_i \leq t < T_i + \sigma_i) \quad (1)$$

where $\mathbf{1}(E)$ is equal to 1 if the event E occurs and to 0 otherwise.

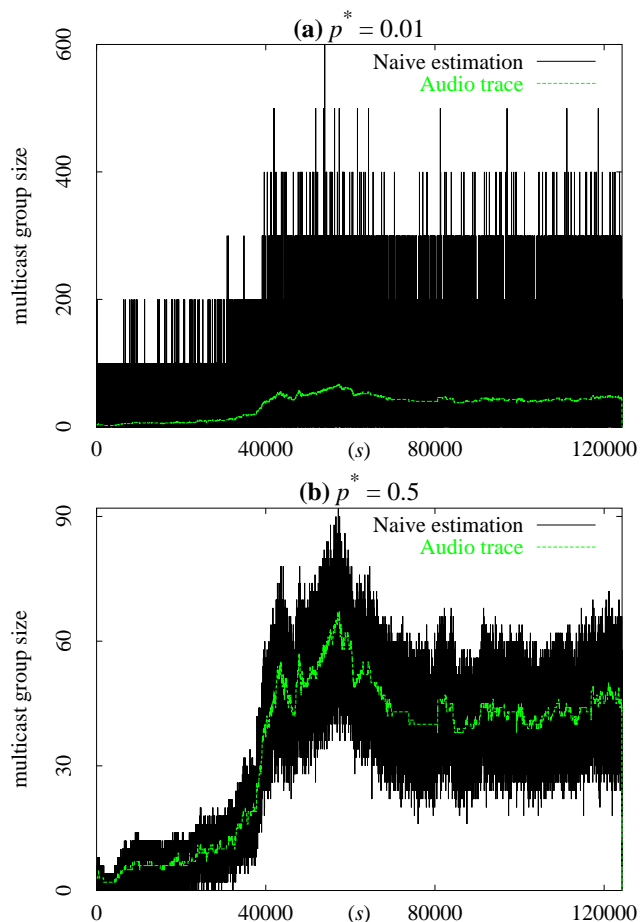


Fig. 1. Evolution of a multicast group size and its estimation using a naive approach ($S = 1$ s).

At times $t = nS, n = 1, 2, \dots$, the source polls all participants (or receivers), with $S > 0$ a given constant. Upon receiving a polling request, a participant sends an ACK back to the source with probability p^* and does not send any feedback information to the source with probability $1 - p^*$. Travel times of polling requests from the source to the receivers and travel times of ACKs from the receivers to the source are supposed to be negligible with respect to S ; we also assume that the processing time needed to generate an ACK at a receiver is negligible with respect to S . Finally, we assume that neither polling requests nor ACKs can be lost. Under these assumptions it is seen that at time nS the source possesses all ACKs sent to it by participants which have been polled at the n -th polling instant (i.e. at time nS). Throughout the paper p^* and S are held fixed.

Given this scheme, our objective is to devise an algorithm for estimating the session size at times $t = nS$ for $n = 1, 2, \dots$

Mainly for mathematical tractability we shall assume from now on that the arrival process is Poisson with rate $\lambda > 0$ and that on-times form a renewal sequence with common exponential distribution with finite mean $1/\mu$, further independent of the arrival process (more general arrival processes can be considered – see Remark II.1). In this setting, the process $\{N(t), t \geq 0\}$ as defined in (1) is nothing but the occupation process in a $M/M/\infty$ queueing system with arrival rate λ and mean service time $1/\mu$ [9]. For such a queue, it is known that

the stationary number of busy servers is distributed according to a Poisson random variable (RV) with parameter $\rho := \lambda/\mu$; in particular, the mean number of busy servers in steady-state is equal to ρ .

Unfortunately, not much is known on the transient behavior of the $M/M/\infty$ for a fixed traffic intensity ρ . We will instead investigate the $M/M/\infty$ queue in heavy traffic. To this end, let us introduce the scaled process $\{N_T(t), t \geq 0\}$ which is identical to the original process $\{N(t), t \geq 0\}$ except for the fact that the arrivals have been speeded up by a factor T , that is, the arrival rate in the $M/M/\infty$ queue is now λT . The mean service rate is kept unchanged and equal to $1/\mu$.

Since $N_T(t) \rightarrow \infty$ a.s. as $T \rightarrow \infty$, we will instead work with the normalized process $\{Z_T(t), t \geq 0\}$ defined by

$$Z_T(t) = \frac{N_T(t) - \rho T}{\sqrt{T}}, \quad t \geq 0. \quad (2)$$

The process $\{Z_T(t), t \geq 0\}$ describes the fluctuations of the scaled process $\{N_T(t), t \geq 0\}$ around its limiting trajectory ρT as $T \rightarrow \infty$.

A nice feature of the process $\{Z_T(t), t \geq 0\}$ is that it converges to a diffusion process as $T \rightarrow \infty$. More precisely, as $T \rightarrow \infty$ the process $\{Z_T(t), t \geq 0\}$ converges in distribution to the Ornstein-Ühlenbeck process $\{X(t), t \geq 0\}$ given by [10, Thm 6.14, p. 155]

$$X(t) = X(0) e^{-\mu t} + \sqrt{2\lambda} \int_0^t e^{-\mu(t-u)} dB(u) \quad (3)$$

where $\{B(t), t \geq 0\}$ is the standard Brownian motion (see also [11, Thm 1, p. 172]). The Ornstein-Ühlenbeck process defined in (3) is an ergodic Markov process and its invariant distribution is a normal distribution with mean zero and variance ρ [12, p. 358].

In the next section we shall devise optimal estimators for the elements of the sequence $\{X(nS), n = 1, 2, \dots\}$ based on Kalman filter theory.

A word on the notation in use: $N(m, v)$ will denote a normal distribution with mean m and variance v and $X \sim N(m, v)$ will denote a RV with distribution $N(m, v)$.

Remark II.1: . The convergence of the process $\{Z_T(t), t \geq 0\}$ to a diffusion process (but with different coefficients than that in (3)) still takes place if the arrival process is replaced by a process slightly more general than a Poisson process [11, Thm 1, pp. 172-173]. On the other hand, if the on-times are generally distributed and the arrival process is Poisson, then one only knows that the process $\{Z_T(t), t \geq 0\}$ converges to a Gaussian process [13].

III. OPTIMAL ESTIMATION

In order to achieve an optimal estimation for the heavy-traffic model, we shall use a Kalman filter that computes the state estimator out of two linear equations: the system dynamic equation and the measurement equation. These two equations are introduced in the next two sections.

Throughout these sections we shall assume that the process $\{X(t), t \geq 0\}$ is in equilibrium at time $t = 0$, namely, $X(0) \sim N(0, \rho)$ (see comment after (3)).

A. System dynamics

From (3) we obtain

$$X(t) = e^{-\mu(t-s)} X(s) + \sqrt{2\lambda} \int_s^t e^{-\mu(t-u)} dB(u)$$

for $0 \leq s \leq t$, from which it follows that

$$\xi_{n+1} = \gamma \xi_n + w_n, \quad n = 0, 1, \dots \quad (4)$$

with $\xi_n := X(nS)$, $\gamma := \exp(-\mu S)$ and

$$w_n := \sqrt{2\lambda} \int_{nS}^{(n+1)S} e^{-\mu((n+1)S-u)} dB(u).$$

The RV's $\{w_n, n = 0, 1, \dots\}$ are i.i.d. with

$$w_n \sim N(0, Q), \quad n = 0, 1, \dots \quad (5)$$

(see e.g. [11, p. 17]) with Q given by

$$\begin{aligned} Q &= 2\lambda \mathbf{E} \left[\int_{nS}^{(n+1)S} e^{-\mu((n+1)S-u)} dB(u) \right]^2 \\ &= 2\lambda \int_{nS}^{(n+1)S} e^{-2\mu((n+1)S-u)} du \\ &= \rho(1 - \gamma^2). \end{aligned}$$

Notice that

$$\xi_n \sim N(0, L), \quad n = 0, 1, \dots \quad (6)$$

with $L := \rho$, under the assumption that $X(0) = \xi_0 \sim N(0, \rho)$, which implies from (4) that

$$\text{cov}(\xi_n, \xi_{n+1}) = \gamma\rho, \quad n = 0, 1, \dots \quad (7)$$

Equation (4) establishes a simple one-order recursive expression relating the state of the limiting process $\{X(t), t \geq 0\}$ between two consecutive polling instants. We shall next derive the corresponding measurement discrete-time equation, which will allow us to use standard optimal estimation techniques.

B. Measurement equation

Let $\zeta^i(n)$ be the indicator function that receiver $i = 1, 2, \dots, N_T(nS)$ has sent an ACK at the n -th polling instant, with $\zeta^i(n) = 1$ if an ACK was sent by receiver i and $\zeta^i(n) = 0$ otherwise. From the definition of the model it is seen that, conditioned on $N_T(nS) = k$, $\zeta^1(n), \dots, \zeta^k(n)$ are i.i.d. Bernoulli RV's with $\mathbf{E}[\zeta^i(n)] = p^*$. The conditional expectation and variance of the number of ACKs $\sum_{i=1}^{N_T(nS)} \zeta^i(n)$ received by the source at time nS are then respectively given by

$$\begin{aligned} \mathbf{E} \left[\sum_{i=1}^{N_T(nS)} \zeta^i(n) \mid N_T(nS) \right] &= N_T(nS) p^* \quad (8) \\ \text{Var} \left[\sum_{i=1}^{N_T(nS)} \zeta^i(n) \mid N_T(nS) \right] &= N_T(nS) p^*(1 - p^*). \quad (9) \end{aligned}$$

We define our normalized measurement equation as

$$Y_T(n) = \frac{\sum_{i=1}^{N_T(nS)} \zeta^i(n) - p^* \rho T}{\sqrt{T}}, \quad n = 0, 1, \dots \quad (10)$$

which can be rewritten as

$$Y_T(n) = p^* Z_T(nS) + V_T(n) \quad (11)$$

with the help of (2), with

$$V_T(n) := \frac{\sum_{i=1}^{N_T(nS)} \zeta^i(n) - N_T(nS)p^*}{\sqrt{T}}. \quad (12)$$

The next step is to let $T \rightarrow \infty$ in (11). It is shown in the Appendix that there exist i.i.d. RV's $\{v_n, n = 0, 1, \dots\}$ with

$$v_n \sim N(0, R), \quad n = 0, 1, \dots \quad (13)$$

with $R := \rho p^* (1 - p^*)$, independent of $\{w_n, n = 0, 1, \dots\}$, such that $\{v_k, k = n, n + 1, \dots\}$ is independent of $\{\xi_k, k = 0, 1, \dots, n\}$ for $n = 0, 1, \dots$, and such that $Y_T(n)$ weakly converges as $T \rightarrow \infty$ to a RV y_n with the representation

$$y_n = p^* \xi_n + v_n, \quad n = 0, 1, \dots \quad (14)$$

The properties enjoyed by the RV's v_n together with (6), (7) and (14) readily imply that

$$y_n \sim N(0, M) \quad (15)$$

with variance $M := \rho p^*$, and

$$\text{cov}(y_n, y_{n+1}) = \gamma \rho (p^*)^2 \quad (16)$$

for $n = 0, 1, \dots$

C. Kalman filter

Equations (4) and (14) represent the equations of a discrete time linear filter, for which we can compute the optimal estimator. Throughout we shall assume that the Gaussian initial condition ξ_0 , the signal noise sequence $\{w_n\}_n$ and the observation noise sequence $\{v_n\}_n$ are all mutually independent.

Let $\hat{\xi}_n$ be an estimator of ξ_n , and denote by $\epsilon_n = \xi_n - \hat{\xi}_n$ the estimation error. The estimator that minimizes the mean square of the estimation error is given by the following Kalman filter (see e.g. [14, p. 347]), which has the following simple recursive structure:

$$P_n = \left((\gamma^2 P_{n-1} + Q)^{-1} + (p^*)^2 / R \right)^{-1} \quad (17)$$

$$K_n = P_n p^* / R \quad (18)$$

$$\hat{\xi}_n = \gamma \hat{\xi}_{n-1} + K_n \left(y_n - p^* \left(\gamma \hat{\xi}_{n-1} \right) \right) \quad (19)$$

for $n = 1, 2, \dots$, with $\hat{\xi}_0 = \mathbf{E}[\xi_0] = 0$ and where the constants γ , R and Q have been defined in Sections III-A and III-B. Equation (17) is called the Riccati equation and P_n gives the variance of the estimation error ϵ_n . In (18) K_n is called the filter gain. Equation (19) is the state estimate equation and it is the sum of an extrapolation term and of an update term.

The above filter minimizes the sum of mean square estimation errors until time nS . One can also (and will from now on) use the stationary version of the Kalman filter, which minimizes the time average mean square estimation error, namely,

$$\begin{aligned} P &= \left((\gamma^2 P + Q)^{-1} + (p^*)^2 / R \right)^{-1} \\ K &= P p^* / R \end{aligned} \quad (20)$$

$$\hat{\xi}_n = \gamma \hat{\xi}_{n-1} + K \left(y_n - p^* \left(\gamma \hat{\xi}_{n-1} \right) \right). \quad (21)$$

P now gives the steady-state variance of the estimation error. It is obtained as the unique positive solution of the algebraic Riccati equation

$$(p^*)^2 \gamma^2 P^2 + (Q (p^*)^2 + R (1 - \gamma^2)) P - R Q = 0.$$

We find

$$\begin{aligned} P &= -\frac{Q (p^*)^2 + R (1 - \gamma^2)}{2 (p^*)^2 \gamma^2} \\ &\quad + \frac{\sqrt{(Q (p^*)^2 + R (1 - \gamma^2))^2 + 4 (p^*)^2 \gamma^2 R Q}}{2 (p^*)^2 \gamma^2}. \end{aligned} \quad (22)$$

We may replace $Q = \rho(1 - \gamma^2)$ and $R = \rho p^*(1 - p^*)$ in (22). Plugging the resulting expression in (20), we can express the filter gain in terms of p^* and γ

$$K = \frac{-(1 - \gamma^2) + \sqrt{(1 - \gamma^2)(1 - \gamma^2(1 - 2p^*)^2)}}{2\gamma^2 p^*(1 - p^*)}. \quad (23)$$

For every n , the error ϵ_n is a normal RV with mean zero and variance P , further independent of the observation y_n [15, p. 240].

D. Membership size estimation

We now return to our original estimation problem, namely, the derivation of an estimate – called \hat{N}_n – for the number of participants $N_T(nS)$ at time nS .

Recall that the process $\{N_T(t), t \geq 0\}$ describes the number of busy servers in an $M/M/\infty$ queue with arrival rate λT and service rate μ . If $N_T(0) = 0$, namely the system is initially empty, then we know by Takács [16, Thm 1, pp.160-161] that $\mathbf{E}[N_T(t)] = \rho T(1 - e^{-\mu t})$ for any time t . In particular, $\mathbf{E}[N_T(t)] = \rho T$ in steady-state (i.e. as $t \rightarrow \infty$).

Motivated by (2), we define \hat{N}_n as follows:

$$\hat{N}_n = \hat{\xi}_n \sqrt{T} + \rho T \quad (24)$$

with $\hat{\xi}_n$ given in (21). We go back to the latter equation. Replacing y_n with $Y_T(n)$ as provided in (10) and using (24), we derive the state estimate equation for \hat{N}_n which reads

$$\begin{aligned} \hat{N}_n &= \gamma(1 - K p^*) \hat{N}_{n-1} + K \sum_{i=1}^{N_T(nS)} \zeta^i(n) \\ &\quad + \rho T(1 - \gamma)(1 - K p^*). \end{aligned} \quad (25)$$

$\sum_{i=1}^{N_T(nS)} \zeta^i(n)$ is the amount of ACKs collected at the n th observation step.

TABLE I
PERCENTILES AND INTERQUARTILE RANGE OF THE RELATIVE ERROR

	Small load $\rho T = 34.1$						Heavy-load $\rho T = 200$					
	25%	50%	75%	90%	95%	Inter	25%	50%	75%	90%	95%	Inter
All exponential	1.9	4.1	7.3	10.6	13.1	5.4	0.8	1.7	2.7	3.8	4.7	1.9
Pareto inter-arrivals	2.6	5.3	9.2	12.9	15.7	6.6	1.5	3.3	5.3	7.2	9.8	3.8
Pareto on-times	4.0	9.2	16.5	26.7	35.7	12.5	3.2	6.1	8.9	12.4	15.1	5.7
All Pareto	3.9	7.8	12.5	18.5	25.6	8.6	1.7	3.5	6.1	9.9	14.5	4.3

Starting from $\mathbf{E}[\hat{\xi}_0] = 0$ it is seen from (21) and (14) that $\mathbf{E}[\hat{\xi}_n] = 0$ for $n = 0, 1, \dots$, which in turn implies from (24) that $\mathbf{E}[\hat{N}_n] = \rho T$. The estimator \hat{N}_n is asymptotically unbiased in the sense that

$$\left| \mathbf{E} \left[\hat{N}_n - N_T(nS) \right] \right| = e^{-\mu n S} \xrightarrow{n} 0.$$

Regarding the variance of the error $e_n := N_T(nS) - \hat{N}_n$, it is a function of the parameter T . We have (use (2)) $e_n = \sqrt{T} \left(Z_T(nS) - \hat{\xi}_n \right)$, so that

$$\text{Var} \left(\frac{e_n}{\sqrt{T}} \right) = \text{Var} \left(Z_T(nS) - \hat{\xi}_n \right).$$

Although we have not been able to prove this result, we believe that $\text{Var} \left(e_n / \sqrt{T} \right) \rightarrow P$ as $T \rightarrow \infty$.

We conclude this section by summarizing the estimation algorithm (ρ and T are known/estimated before hand)

Initialization:

Set p^* and S to the desired values, and \hat{N}_0 to ρT ($\hat{\xi}_0 = 0$). Compute once for all $\gamma = \exp(-\mu S)$ and the filter gain K according to (23).

n th observation step:

Collect the ACKs received in the interval of time $((n-1)S, nS]$ and compute \hat{N}_n as in (25).

IV. SIMULATIONS

Our estimator has been derived under a set of various statistical assumptions that may be violated in practice (Poisson arrivals, exponential on-times, heavy-traffic regime). In this section, we investigate the robustness of our estimator and try to identify situations where it works well/poorly. To do so, we have conducted various simulations (using a *C* program) where some or all of the assumptions needed for the derivation of the estimator are violated.

Four types of simulations have been performed. For each simulation, the parameters λ and T are taken to be equal to 1 and $1/185.9s^{-1}$, respectively, and the run time is 124240s. These values have been measured on a real trace (see Section V for details on the traces we have used). The ACK probability p^* and the ACK interval S have been set to 0.01 and 1.0s, respectively (see comments in Section VI).

Two figures are associated with each simulation depending on the load (defined as ρT) of the system: $\rho T = 34.1$ referred to as “small load” (left-hand side figures) and $\rho T = 200$ referred to as “heavy-load” (right-hand side figures). Each figure

displays three curves: the simulated data, the estimated data and the mean load ρT . For each simulation, the performance of the estimator is collected in Table I. The first column gives the 25th percentile, the second the median value, etc. The last column reports the interquartile range² of the relative error.

In the first simulation the users join the multicast group according to a Poisson process and their on-times are exponentially distributed. The Poisson assumption for the joining process is fairly realistic as mentioned in [17]. The validity of the exponential assumption for the on-times has been observed for short sessions. The obtained results are reported in Figures 2(a). Both for small and heavy loads the estimated value appears to be very close to the true value; in particular, the relative error defined as $\frac{|N_T(nS) - \hat{N}_n|}{N_T(nS)}$, is less than 13.1% (resp. less than 4.7) most of the time when $\rho T = 34.1$ (resp. $\rho T = 200$) (see Table I for details).

In the second simulation the on-times are still exponentially distributed but now the inter-arrivals are Pareto distributed with shape parameter α equal to 1.1, leading to an infinite variance of the inter-arrival times. The results are displayed in Fig. 2(b). We first observe that both the estimator and the real values are far away from the “limiting trajectory” when the load is small. This is due to the infinite variance of the inter-arrival times which prevents the stationary regime to be reached rapidly. Nevertheless, the accuracy of the estimator is still remarkable both at small and heavy loads, as can be seen in Table I.

In the third simulation the arrival process is Poisson and the on-times are Pareto distributed with shape parameter α equal to 1.1. Both the estimator and the real values are far away from the limiting trajectory for both loads. When the load is small, the accuracy of the estimator is not as good as in the previous simulations but it is still fair. We suspect that this lack of accuracy is more a consequence of the small measured load (18.1) than of the Pareto on-time assumption. The accuracy of the estimator dramatically increases as the load increases. See Fig. 2(c) and Table I for details.

In the fourth and last simulation all assumptions are simultaneously violated: both the inter-arrival times and the on-times are Pareto distributed with shape parameter α equal to 1.1. The overall performance of the estimator is better than in the third simulation. See Fig. 2(d) and Table I for details. Table II contains the mean and variance of the error e_n . The expected mean is 0 and the (conjectured) expected variance is TP (see Section

²The interquartile range is the difference between the upper (75th percentile) and lower (25th percentile) quartiles of the data sample (upper and lower bounds of the center half of the data values). It describes the relative concentration of the data around the median value.

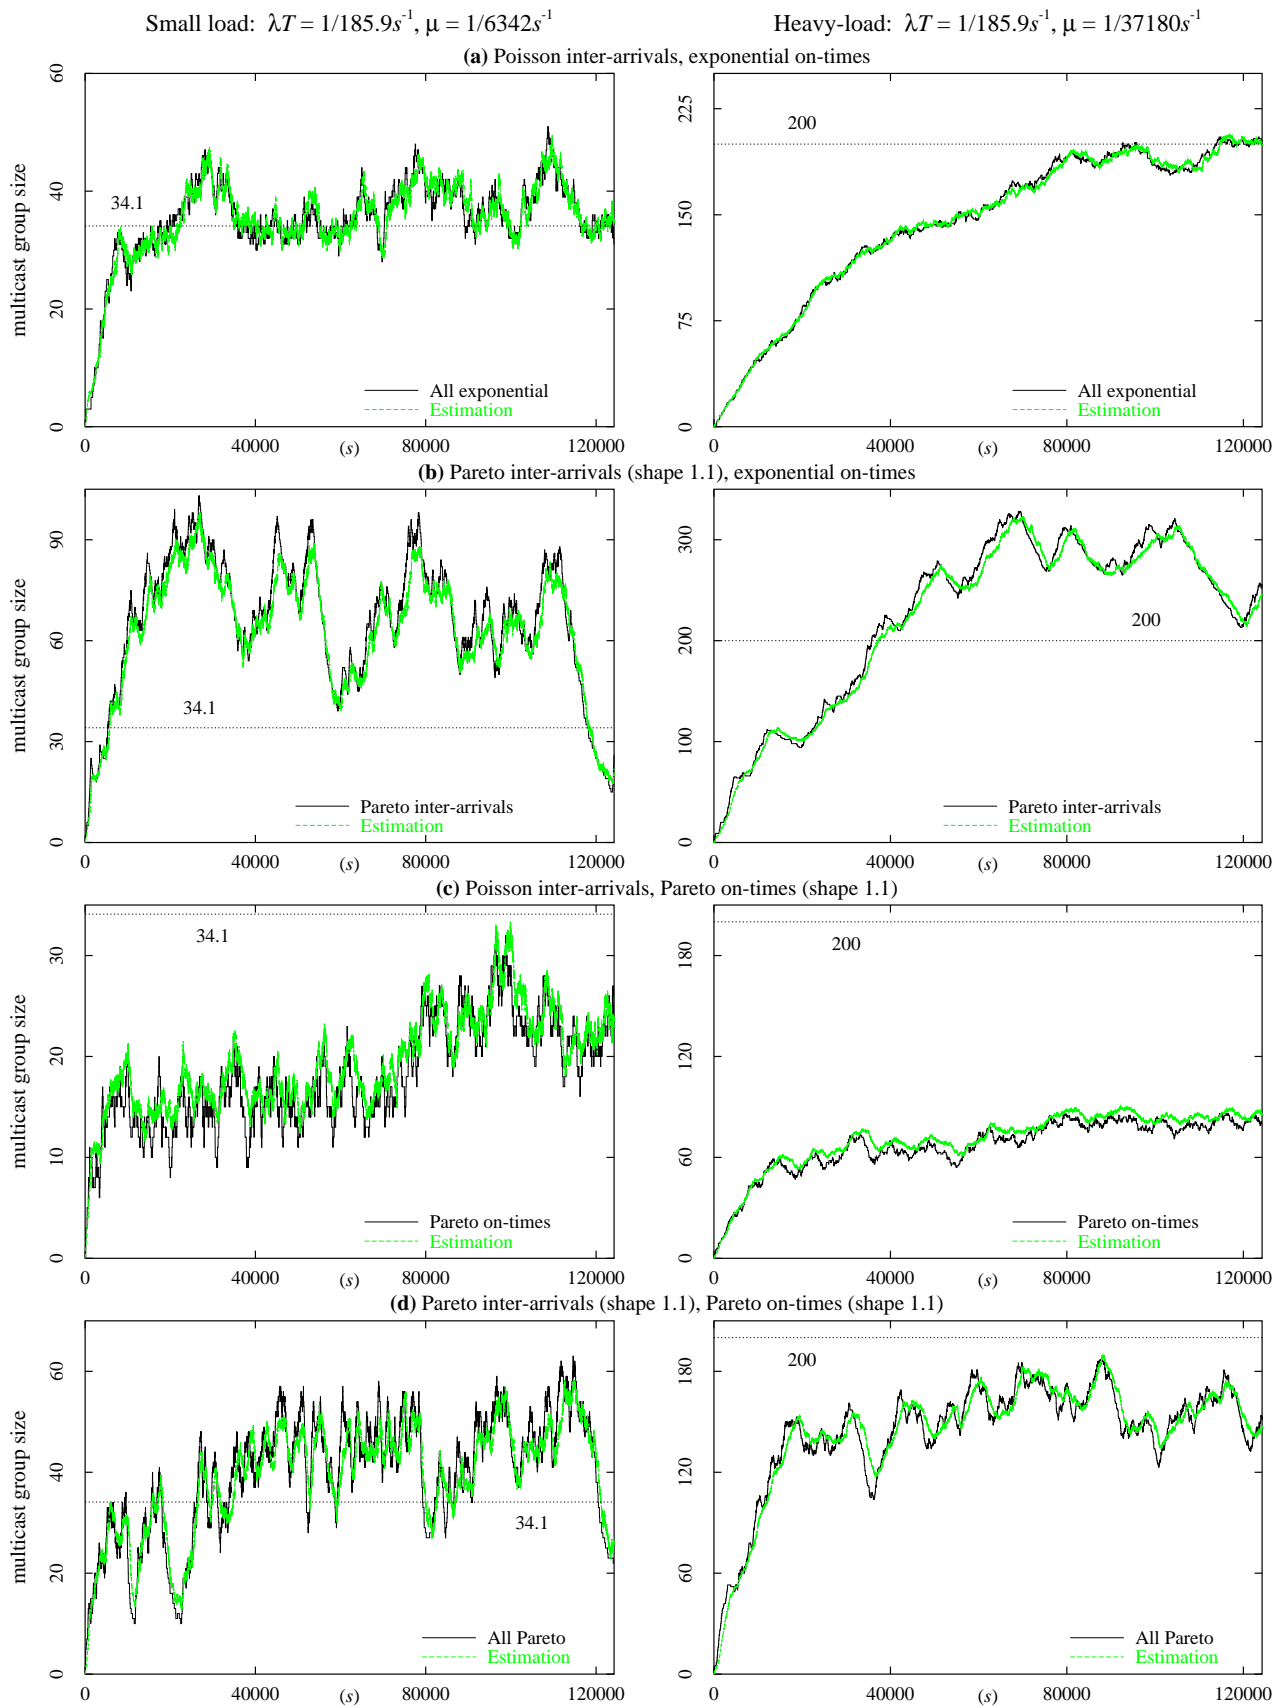


Fig. 2. Estimation of the multicast group size over time: $p^* = 0.01$, $S = 1s$.

TABLE II
MEAN AND VARIANCE OF THE ERROR e_n

	$\rho T = 34.1$		$\rho T = 200$	
	Mean	Var	Mean	Var
All exponential	0.07	4.7	-0.006	9.95
Pareto inter-arrivals	2.8	19.7	3.6	67.3
Pareto on-times	-1.2	4.1	-3.8	10.1
All Pareto	0.6	16.9	-0.4	57.9

III-D). “Mean” denotes the measured average and “Var” denotes the measured variance. Looking at Table II, we can easily observe that

- There is a small bias when the distributions of the inter-arrival times and of the on-times are different, that is, when one of them is Pareto and the other one is exponential (c.f. second and third lines). The negative bias -3.8 in the heavy-load case can be observed in Fig. 2(c) where the estimation is clearly above the simulated group size;
- The variances measured when the arrival process is Poisson are very close to each other for both values of the workload (see 1st and 3rd lines) and they are not too far from the expected variances;
- The variances measured when the inter-arrival times are Pareto distributed are both far from the expected variances but are relatively close to each other (see 2nd and 4th lines).

V. VALIDATION WITH REAL TRACES

An extensive study of the characterization of MBone session dynamics is due to Almeroth and Ammar [17], [18]. They have developed a tool called *Mlisten*³ that collects the join/leave times for multicast group members in MBone sessions. We have applied our algorithm to some of these traces⁴ collected in 1996. We were not able to analyze more recent traces. As stated in [18] (see Section V-A and Fig. 3), the joining process is reasonably close to a Poisson process. As to the on-times two cases have to be distinguished depending on the duration of a session. For long sessions some people will join for very long periods while others will join only for a few minutes. In this case, the Zipf distribution fits well in the collected data. In the case where sessions are short, the maximum membership duration is much shorter than for long sessions, thereby eliminating long on-times.

We have run our algorithm on two different traces, one collected from a short audio session that started on 9th of December 96 and lasted for 1 day 10 hours 30 minutes and 40 seconds, i.e. 124240 seconds; the other one results from a long audio session that lasted from 18th of November 96 to 10th of December 96 (21 days 12 hours 37 minutes and 27 seconds, that is 1859847 seconds). Fig. 3 plots the actual group size and its estimation for each session.

The values of the arrival rate T (throughout $\lambda = 1$ so that the arrival rate is T) and the expected on-times $1/\mu$ were extracted

from the traces. For the short duration session, the measured “load” is $\rho T = 34.1$; for the long duration session, the measured load is $\rho T = 63.5$.

We have observed (see Fig. 3(a)) that our estimator does not work well case the session gathers a few participants. In this case, it overestimates the size of the group; the absolute error $|N_T(nS) - \hat{N}_n|$ is not significant but the relative error is very high and approaches 100% (see around 2 a.m. in Fig. 3(a)). This behavior was already reported in Section IV. The percentiles and the interquartile range of the relative error for both sessions are listed in Table III.

TABLE III
PERCENTILES AND INTERQUARTILE RANGE OF THE RELATIVE ERROR

ρT	25%	50%	75%	90%	95%	Inter
34.1	2.3	5.2	11.8	44.9	61.2	9.5
63.5	1.3	2.7	4.7	7.7	13.1	3.3
38.0	1.8	4.2	7.3	11.4	15.9	5.4

The data set corresponding to the short audio session exhibits two very different parts: very few users are connected during the first quarter of the data set, thereby suggesting that this part was recorded before the start of the transmission. The remaining part of the data set records participants activity during the transmission.

Our algorithm has first been run on the entire trace. As expected, it performs poorly on the first quarter of the trace; its performance improves dramatically on the remaining part of the trace.

The first quarter of the data set has then been removed and our algorithm has been run on the resulting trace (we found $T = 1/157.7s^{-1}$, $\mu = 1/5994.6s^{-1}$ and $\rho T = 38$). Most of the time, the relative error is less than 15.9% (see Table III), which is a satisfactory result.

Table IV reports the mean and variance of the error e_n . As before, “Mean” is the measured average, “Var” is the measured variance and “ TP ” is the (conjectured) expected variance. We can see from this table that, as expected, the measured average is close to 0. Also notice that the measured variance and the expected variance are close to each other.

TABLE IV
MEAN AND VARIANCE OF THE ERROR e_n

ρT	Mean	Var	TP
34.1	0.219	6.070	5.514
63.5	0.017	8.083	6.723
38.0	0.338	6.339	6.303

A. Distribution fits

The distribution fits for the inter-arrival time process and the on-time process for both the short and long duration sessions are presented in Figs. 3(a) and 3(b). We have found that for both types of sessions, the joining process is not Poisson and on-times are not exponentially distributed.

³ Available at <http://www.cc.gatech.edu/computing/Telecomm/mbone/>

⁴ Available at <ftp://ftp.cc.gatech.edu/people/kevin/release-data/>

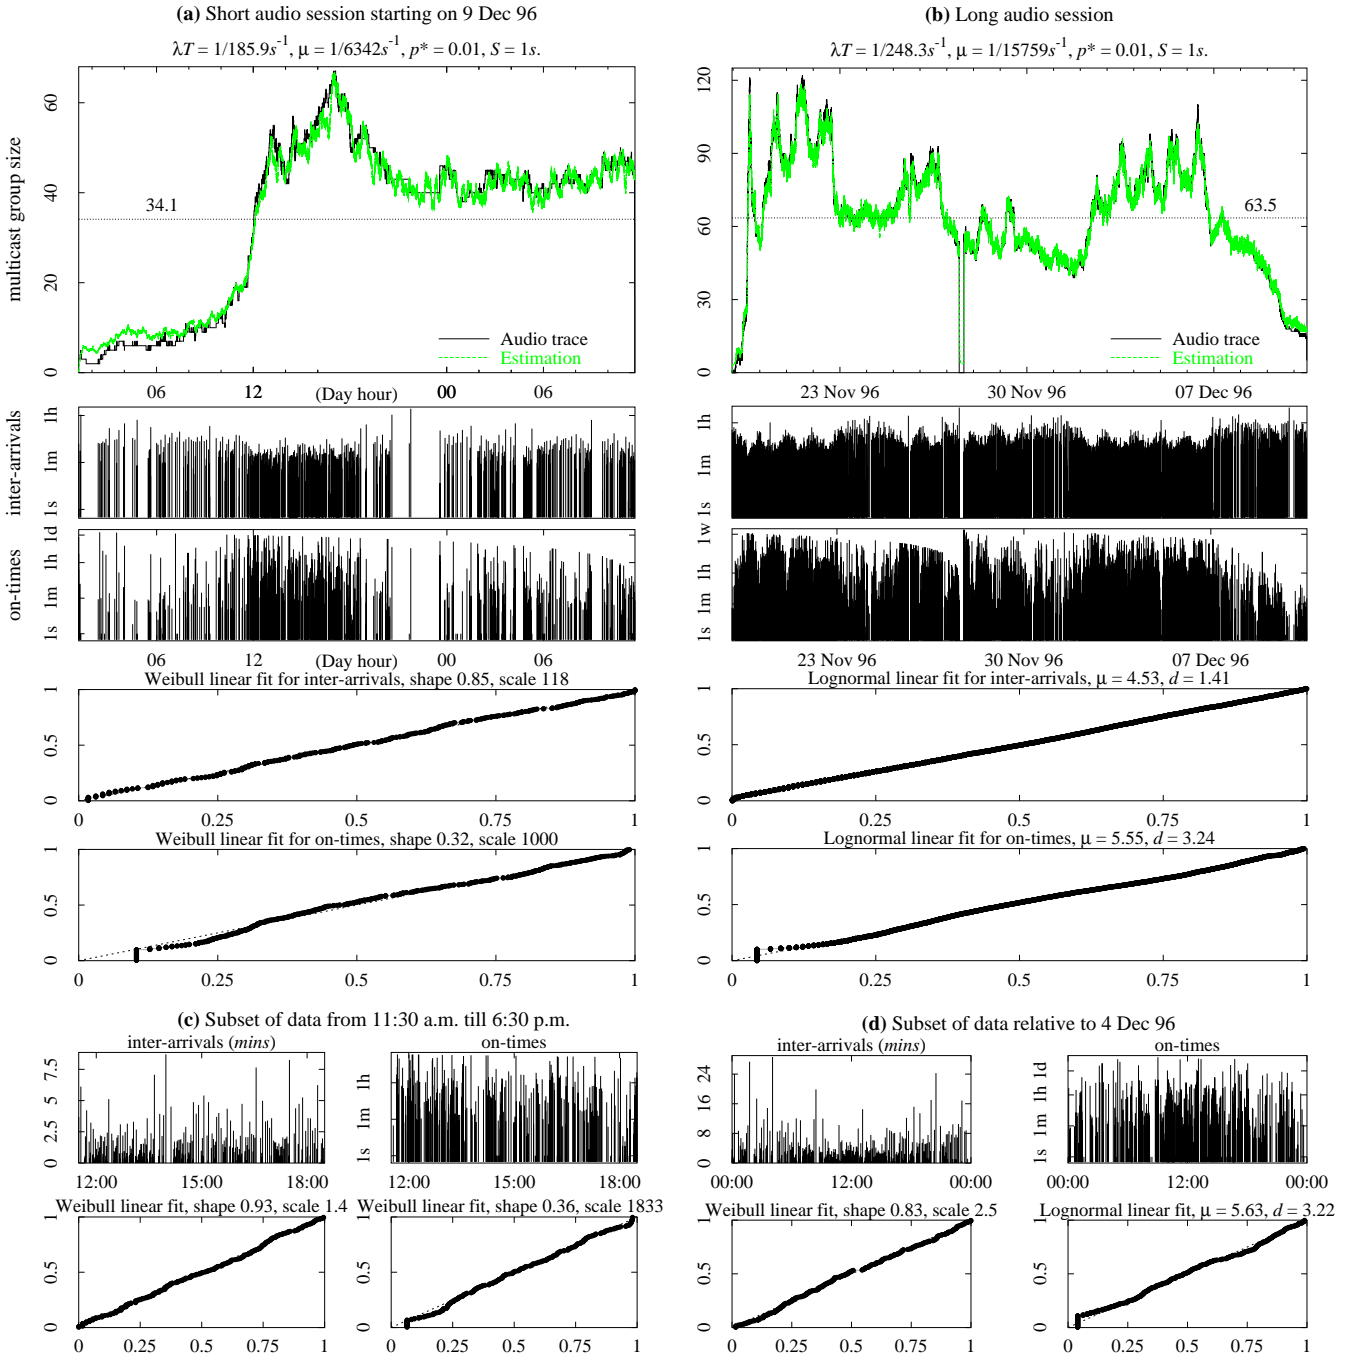


Fig. 3. Estimation of the multicast group size using audio traces and probability plots for the observed data.

For the entire short audio session, the inter-arrival time distribution is well represented by a Weibull distribution with shape parameter 0.85 and the on-time distribution fits well a Weibull distribution with shape parameter 0.32. Recall that the smaller the shape parameter the heavier the tail (a Weibull distribution with shape parameter 1 is an exponential distribution). Our results are different from the results reported in [18] which is explained by the fact that they have only analyzed the part of the trace with the highest human activity. See Fig. 3(a).

For the entire long audio session, the inter-arrival time distribution is well represented by a Lognormal distribution with parameters $\mu = 4.53$ and $d = 1.41$ and the on-time distribution

fits well a Lognormal distribution with parameters $\mu = 5.55$ and $d = 3.24$. Recall that the Lognormal distribution is long-tailed; the higher the parameter d , the longer the tail. See Fig. 3(b).

To conclude this section, we would like to point out that our estimator seems to be very robust to changes in the distribution laws. Although the assumptions that the inter-arrival times and the on-times have an exponential distribution is crucial in the theory that we have developed in Sections II and III, it is interesting to note that our estimator still performs well for other distributions, including various subexponential distributions (Pareto, Lognormal, Weibull).

VI. OPEN ISSUES

Our main concern is the choice of the couple (p^*, S) . Until now we have used $p^* = 0.01$ and $S = 1s$ in all experiments. These values were chosen in order to have a good estimation. Empirically, we have seen that the estimation is also good if both p^* and S are increased/decreased. If we only increase S , the ACK process will exhibit less correlations, and the estimation will deteriorate. In this case, we will need to increase p^* accordingly in order to enhance the estimation. On the other hand, if we decrease p^* then the source will collect less feedback information and S should be decreased in order to increase the correlation of the ACK process.

From the network point of view, the best would be to have small p^* and large S so as to minimize congestion. From the estimation point of view, the best would be to have small S and large p^* so as to enhance the quality of the estimator. Therefore, a trade-off has to be found based on the group size, the available bandwidth of the underlying network and the capacity of the source (or of any receiver who wants to dynamically estimate the group size). We can think of three criteria to choose (p^*, S) in order to

- (i) Avoid feedback implosion
- (ii) Allow for a maximal volume of ACKs
- (iii) Insure a predefined estimation quality.

The first criterion is the most constraining as it limits the amount of ACKs to be generated at each ACK epoch. The expected number of ACKs to be generated at time nS is $p^*N_T(nS)$ which is estimated as $p^*\hat{N}_n$. This criterion is expressed as $p^*\hat{N}_n \leq \alpha$ which means that the highest expected number of ACKs during S seconds is equal to α . When applying our algorithm to the real traces, we obtained up to 6 ACKs per second ($S = 1$) for the short session and up to 8 ACKs per second for the long session. If this exceeds the desired threshold, it means that $p^*\hat{N}_n > \alpha$ and we should then decrease p^* slightly, which could deteriorate the quality of the estimation.

Regarding the second criterion, it is less constraining in the sense that it allows for higher p^* . This criterion takes advantage of the fact that receivers are not synchronized and the generated ACKs are dispersed over S seconds. So we may allow for a larger volume of ACKs to be generated as long as it does not exceed some threshold during some I seconds. This is expressed as $p^*\hat{N}_n \frac{I}{S} \leq \alpha$ where I is in the order of μs or even ms . Increasing I for the same α would be more constraining and letting $I = S$ gives us the first criterion. For $I > S$, this criterion becomes the most constraining one.

For overprovisioned networks, the only worthwhile criterion is to insure some estimation quality. Fig. 4 plots the expected variance of the error e_n over $\gamma = \exp(-\mu S)$ for different values of p^* . Two couple of values were retained for λT and μ , taken from the couple of traces studied in Section V. In each plot the horizontal line corresponds to $p^* = 0.01$ and $\gamma = \exp(-\mu)$ ($S = 1s$). It gives the expected variance of the error obtained in the corresponding real traces. Observe that several couples (p^*, γ) lead to the same variance, but while the r.h.s. values correspond to high correlations between the ACKs, the l.h.s. values result in independent ACKs. For high γ (small S) the estimation reproduces the dynamics of the group size, while for small γ (large S) it reproduces the mean value of the group size

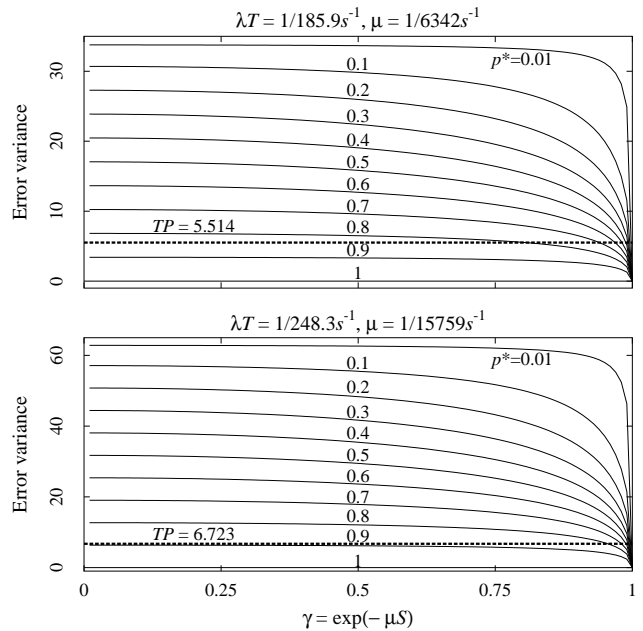


Fig. 4. Expected error variance for different values of the ACK probability.

(load). However, the variance of the error would be the same in both cases if we consider appropriate values for p^* . We are interested in the r.h.s. values and we may satisfy one of the previous criteria in order to choose the parameters values.

VII. CONCLUSION

In this paper, we have developed a robust estimator for computing the size of a multicast group. Our scheme allows the user to track on-line the variation of the audience over time. We have modeled the group dynamics as an $M/M/\infty$ queue and established our results under the assumption that this queue is in heavy-traffic. In this regime the backlog process of the $M/M/\infty$ queue is “close” to a diffusion process that can be used to cast our estimation problem into the appealing framework of Kalman filter theory. Using this theory we have derived an estimator that minimizes the variance of the error. We have carried out several simulations to test the robustness of our estimator in the case where the arrivals are not Poisson and/or the on-times are not exponentially distributed. The estimator has also been computed on real audio traces and its performance have been shown to be excellent. Ongoing research include an optimal selection of (p^*, S) with respect to some cost criterion (e.g. expected number of ACKs bounded by a given constant) as well as extensions to the case where both p^* and S have to be estimated, the objective being to better react to sudden changes in the multicast population.

Acknowledgment: The authors wish to thank Prof. O. Zeitouni for helpful suggestions.

APPENDIX

Define

$$Z(m, n) := \frac{\sum_{i=1}^m \zeta_i(n) - mp^*}{\sqrt{m}}.$$

Observe that

$$V_T(n) = Z(N_T(nS), n) \sqrt{\frac{N_T(nS)}{T}} \quad (26)$$

where $V_T(n)$ is defined in (12)).

The RV $Z(m, n)$ converges weakly as $T \rightarrow \infty$ to a normal RV ℓ_n with mean zero and variance $p^*(1-p^*)$. Equivalently, for any bounded continuous function f ,

$$\lim_{m \rightarrow \infty} \mathbf{E}[f(Z(m, n))] = \mathbf{E}[f(\ell_n)].$$

Since $N_T(nS)/T$ converges P a.s. to ρ as $T \rightarrow \infty$ [10, Thm 6.13, pp. 153-154], it follows that

$$\lim_{T \rightarrow \infty} \mathbf{E}[f(V_T(n)) | N_T(nS)] = \mathbf{E}[f(v_n)] \quad (27)$$

where $v_n = \rho \ell_n$ is a normal RV with mean zero and variance $\rho p^*(1-p^*)$.

Let f and g be two arbitrary bounded continuous functions. Then

$$\begin{aligned} \lim_{T \rightarrow \infty} \mathbf{E}[f(V_T(n)) g(Z_T(nS))] &= \\ &= \lim_{T \rightarrow \infty} \mathbf{E}[\mathbf{E}[f(V_T(n)) g(Z_T(nS)) | N_T(nS)]] \\ &= \lim_{T \rightarrow \infty} \mathbf{E}[\mathbf{E}[f(V_T(n)) | N_T(nS)] g(Z_T(nS))] \\ &= \lim_{T \rightarrow \infty} \mathbf{E}[\{\mathbf{E}[f(V_T(n)) | N_T(nS)] - f(v_n)\} \\ &\quad \times \{g(Z_T(nS)) - g(\xi_n)\} + \mathbf{E}[f(v_n)] \mathbf{E}[g(\xi_n)]] \\ &= \mathbf{E}[f(v_n)] \mathbf{E}[g(\xi_n)] \end{aligned} \quad (28)$$

$$= \mathbf{E}[f(v_n)] \mathbf{E}[g(\xi_n)] \quad (29)$$

where (28) follows from (27) together with the weak convergence of $Z_T(nS)$ to ξ_n as $T \rightarrow \infty$, and where the last equality follows from the bounded convergence theorem.

On the probability space that carries the RV's $\{\xi_n, w_n\}_n$ one can always construct the RV's v_n so that they are i.i.d. with a normal distribution with mean zero and variance $\rho p^*(1-p^*)$, further independent of $\{w_n, n \geq 0\}$ and such that, for every $n \geq 0$, the $\{v_k, k \geq n\}$ are independent of $\{\xi_0, \xi_1, \dots, \xi_n\}$. Under this construction, we deduce from (29) that

$$\lim_{T \rightarrow \infty} \mathbf{E}[f(V_T(n)) g(Z_T(nS))] = \mathbf{E}[f(v_n) g(\xi_n)]$$

or, equivalently, that $(Z_T(nS), V_T(n))$ converges weakly to (ξ_n, v_n) as $T \rightarrow \infty$ (Hint: choose $f(x) = \exp(it_1 x)$ and $g(x) = \exp(it_2 x)$ with t_1 and t_2 real numbers). We deduce from this result that $Y_T(n)$ defined in (10) converges weakly as $T \rightarrow \infty$ to a RV y_n such that $y_n = p^* \xi_n + v_n$.

REFERENCES

- [1] H. Schulzrinne, S. Casner, R. Frederick, and V. Jacobson, "RTP: a transport protocol for real-time applications," RFC 1889, Network Working Group, January 1996.
- [2] J. Rosenberg and H. Schulzrinne, "Timer reconsideration for enhanced RTP scalability," in *Proc. of IEEE INFOCOM'98, San Francisco, CA USA*, March/April 1998, vol. 1, pp. 233-241.
- [3] S. Floyd, V. Jacobson, S. McCanne, C. Liu, and L. Zhang, "A reliable multicast framework for light-weight sessions and application level framing," in *Proc. of ACM SIGCOMM'95, New York City, NY USA*, 1995, pp. 342-356.

- [4] J.-C. Bolot, T. Turletti, and I. Wakeman, "Scalable feedback control for multicast video distribution in the Internet," in *Proc. of ACM SIGCOMM'94, London, UK*, September 1994, pp. 58-67.
- [5] J. Nonnenmacher and E. Biersack, "Optimal multicast feedback," in *Proc. of IEEE INFOCOM'98, San Francisco, CA USA*, March 1998, vol. 3, pp. 964-971.
- [6] J. Nonnenmacher and E. Biersack, "Scalable feedback for large groups," *IEEE/ACM Trans. on Networking*, vol. 7, no. 3, pp. 375-386, June 1999.
- [7] T. Friedman and D. Towsley, "Multicast session membership size estimation," in *Proc. of IEEE INFOCOM'99, New York City, NY USA*, March 1999, vol. 2, pp. 965-972.
- [8] C. Liu and J. Nonnenmacher, "Broadcast audience estimation," in *Proc. of IEEE INFOCOM 2000, Tel Aviv, Israel*, March 2000, vol. 2, pp. 952-960.
- [9] L. Kleinrock, *Queueing systems*, vol. 1, John Wiley & Sons, 1975.
- [10] P. Robert, *Réseaux et fi les d'attente: méthodes probabilistes*, Springer, 2000.
- [11] A. A. Borovkov, *Asymptotic methods in queueing theory*, John Wiley & Sons, 1984.
- [12] I. Karatzas and S. E. Shreve, *Brownian motion and stochastic calculus*, Springer, 1991.
- [13] W. A. Rosenkrantz and J. Horowitz, "Statistical analysis of variance-time plots used to estimate parameters of a long-range dependent process," Submitted for publication, February 2001.
- [14] R. F. Stengel, *Stochastic optimal control, theory and application*, John Wiley & Sons, 1986.
- [15] P. Whittle, *Optimal Control. Basics and Beyond*, John Wiley & Sons, 1996.
- [16] L. Takács, *Introduction to the theory of queues*, Oxford University Press, Inc., 1962.
- [17] K. Almeroth and M. Ammar, "Multicast group behavior in the Internet's Multicast Backbone (MBone)," *IEEE Communications magazine*, vol. 35, pp. 224-229, June 1997.
- [18] K. Almeroth and M. Ammar, "Collecting and modeling of the join/leave behavior of multicast group members in the MBone," in *Proc. of HPDC'96, Syracuse, NY USA*, August 1996, pp. 209-216.