



Estimating Membership in a Multicast Session

Sara Alouf, Eitan Altman, Chadi Barakat, Philippe Nain

► To cite this version:

Sara Alouf, Eitan Altman, Chadi Barakat, Philippe Nain. Estimating Membership in a Multicast Session. RR-4391, INRIA. 2002. inria-00072197

HAL Id: inria-00072197

<https://inria.hal.science/inria-00072197>

Submitted on 23 May 2006

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Estimating Membership in a Multicast Session

Sara Alouf — Eitan Altman — Chadi Barakat — Philippe Nain

N° 4391

Février 2002

THÈME 1



*rapport
de recherche*

Estimating Membership in a Multicast Session

Sara Alouf , Eitan Altman , Chadi Barakat* , Philippe Nain

Thème 1 — Réseaux et systèmes
Projet Mistral

Rapport de recherche n° 4391 — Février 2002 — 21 pages

Abstract: We propose two novel on-line estimation algorithms to determine the size of a dynamic multicast group. We first use a Wiener filter to derive an optimal estimator for the membership size of the session in case the join process is Poisson and the lifetime of participants is distributed exponentially. We next develop the best first-order linear filter from which we derive an estimator that holds for any lifetime distribution. We apply this approach to the case where the lifetime distribution is hyperexponential. Both estimators hold under any traffic regime. Applying both estimators on real video traces, we find that both schemes behave well, one of which performs slightly better than the other in some cases. We further provide guidelines on how to tune the parameters involved in both schemes in order to achieve high quality estimation while simultaneously avoiding feedback implosion.

Key-words: on-line estimation, multicast, $M/G/\infty$ queue, Wiener filter, validation.

* This work was carried out while this author was holding a post-doc position at ICA - DSC - EPFL.

Estimation de la taille d'une population multipoint

Résumé : Ce rapport propose deux nouvelles techniques pour l'estimation en ligne de la taille d'une population multipoint dynamique. La première de ces techniques se base sur la théorie de Wiener sur le filtrage linéaire pour donner un estimateur optimal dans le cas d'arrivées Poissonniennes et de temps de connexions exponentiellement distribués. La deuxième approche consiste à développer le filtre linéaire optimal d'ordre 1 qui retourne un estimateur valable pour des temps de connexions de loi générale. Cette technique est illustrée dans le cas où les temps de connexions sont de distribution hyperexponentielle. Les deux estimateurs sont testés sur des traces vidéo et les estimations obtenues sont de bonne qualité. En vue de contrôler le volume d'acquittements généré, tout en fournissant une bonne qualité d'estimation, quelques recommandations sont fournies sur le choix des paramètres en jeu.

Mots-clés : estimation en ligne, multipoint, file d'attente $M/G/\infty$, filtre de Wiener, validation.

1 Introduction

This paper addresses the estimation of the size of a dynamic multicast session. Estimating the size of a multicast session can be quite useful to many protocols like RTP [15], [14] and SRM [6] requiring such estimates for feedback suppression, as well as to several applications. For instance, Bolot, Turetti 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 [12], [13] 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], [12], [13] which have point estimators and further add some contributions to each. Another timer-based feedback scheme is proposed in [11] 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.

In a recent paper [3], the authors consider a polling scheme where each participant sends an acknowledgement (ACK) every S seconds with probability p and does not send anything with probability $1 - p$. They use a diffusion approximation for the heavy-traffic regime. Under the assumptions of Poisson arrival process of participants and exponentially distributed lifetimes, the diffusion approximation yields linear dynamics which allows the authors to design the optimal filter using Kalman filter theory.

In this paper we consider the same polling scheme as in [3] and propose two novel algorithms for estimating the membership. Our first approach is based on a Wiener filter, which provides the optimal dynamic estimator among all linear estimators. The dynamics is not required to be linear unlike in the case of the Kalman filter, which allows us to get rid of the heavy-traffic assumption made in [3]. Yet, in order to obtain explicit expressions for the parameters of the Wiener filter, we still have to assume that participants join the session according to a Poisson process, and that the time during which they stay in the multicast session, hereafter referred to as *on-time*, has an exponential distribution. Under these assumptions we design the optimal linear estimation scheme that turns out to require a filter of order one. Motivated by this structure, we then design an efficient estimation scheme

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

for generally distributed on-times. To that end, we identify the optimal filter among all linear filters of order one. We illustrate this approach in the case of hyperexponentially distributed on-times. Unlike the estimator developed with the Wiener filter, the latter estimator is valid under any traffic regime. Both estimators are then tested on real video traces. Despite the fact that these traces violate the assumptions under which the estimators have been derived, very good performance are observed.

The paper is organized as follows: the mathematical model of the membership is introduced in Section 2. The theory of Wiener filters is briefly presented in Section 3 and its application to the $M/M/\infty$ model comes in Section 4. The optimal first-order linear filter is developed in Section 5. Section 6 proposes some guidelines on how to choose parameters p and S . The robustness of both estimators is addressed through validations on real video traces in Section 7 and concluding remarks are given in Section 8.

2 Multicast membership modeled as an $M/G/\infty$ queue

We consider a multicast group that participants join and leave at random times. Let T_i and $T_i + D_i$ be the join time and leave time, respectively, of the i -th participant. In the following, D_i is called the on-time of the i -th participant. Let $N(t)$ be the number of participants at time t or, equivalently, the size of the multicast session at time t . Clearly

$$N(t) = \sum_{i \geq 1} \mathbf{1}(T_i \leq t < T_i + D_i) \quad (1)$$

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

We shall assume that the join times form a Poisson process (with constant intensity $0 < \lambda = 1/\mathbf{E}[T_{i+1} - T_i]$) and that the on-times form a renewal sequence of random variables (rvs) with common probability distribution $\Psi(x) = P(D_i < x)$, $0 < \mathbf{E}[D_i] < \infty$, further independent of the join times. In the following D will denote a generic rv with probability distribution $\Psi(x)$.

In the queueing terminology, $\{N(t), t \geq 0\}$ represents the occupation process (number of busy servers) in an $M/G/\infty$ queue [10].

At times $t = nS$, $n = 0, 1, \dots$, with $S > 0$ a constant, each participant to the multicast session sends an ACK to the source with probability $0 < p < 1$ and does not send any feedback information to the source with probability $1 - p$. Note that in practice the source will have to regularly multicast the pair (p, S) to ensure that each participant will know these values. We will assume that S is large enough (of the order of the second) so that the travel times of ACKs from participants to the source and the processing time needed to generate an ACK can be neglected. Finally, we assume that ACKs cannot be lost (this assumption can be relaxed if one knows the loss probability). Hence, all ACKs sent by participants at time nS are available at once at the source. Throughout the paper p and S are held fixed (see Section 6 for possible extensions).

Let Y_n be the number of acks received by the source at time nS . Based on the knowledge of Y_1, \dots, Y_n our objective is to find an optimal estimator (in a sense to be defined below)

\hat{N}_n for $N_n := N(nS)$, the size of the multicast group at time nS . In the filtering parlance, Y_n is an input signal and we want to generate another signal \hat{N}_n that is as close as possible to an unknown signal N_n , in the sense of minimizing the mean square error.

For later use we briefly review some results on the $M/G/\infty$ queue. In steady-state, the number N of busy servers is a Poisson random variable with parameter $\rho := \lambda E[D]$, namely, $P[N = j] = \rho^j \exp(-\rho)/j!$. In particular, both the mean and the variance of the number of busy servers are equal to ρ . The autocorrelation function of the stationary version of the process $\{N(t), t \geq 0\}$, also denoted by $\{N(t), t \geq 0\}$, is given by [5, Eqn (5.39)]

$$\text{Cov}(N(t), N(t+h)) = \lambda \int_{|h|}^{\infty} P(D > u) du. \quad (2)$$

In the following we will denote by $\text{Cov}_X(\cdot)$ the autocorrelation function of any second-order discrete-time stationary process $\{X_n, n = 0, 1, \dots\}$. With this notation and the definition of the process $\{N_n\}_n$, we see from (2) that

$$\text{Cov}_N(k) = \rho \gamma^{|k|}, \quad k = 0, \pm 1, \dots, \quad (3)$$

with $\gamma := \exp(-\mu S)$, when the on-times $\{D_i\}_i$ are exponentially distributed with mean $1/\mu$.

Throughout the paper we will assume that

$$\sum_{k \geq 0} \text{Cov}_N(k) < \infty. \quad (4)$$

In other words, we will exclude the situation where the on-times are heavy-tailed (e.g. Pareto distribution).

3 Wiener filter

Our objective is to transform a signal Y_n (noisy observation) into another signal \hat{N}_n (estimator) that is the closest to an unknown signal N_n . By closest we mean that the mean error is zero (i.e. $\mathbf{E}[\hat{N}_n] = \mathbf{E}[N_n]$) and that the mean square error is minimized.

Such a transformation can be achieved by the Wiener filter that identifies the optimal linear filter [8]. This approach gives the transfer function of the linear filter, which can be transformed back to the time domain to obtain the impulse response of the filter. From the impulse response of the filter, the expression of \hat{N}_n as a function of Y_n and, possibly, of $\hat{N}_{n-1}, \hat{N}_{n-2}, \dots, \hat{N}_0$, can be found. We will detail this procedure below.

Since a filter that minimizes the mean square error when the underlying processes are centered also minimizes the mean square error when the same processes are non-centered, we will derive the Wiener filter for the centered (stationary) versions of processes $\{N_n\}_n$, $\{\hat{N}_n\}_n$ and $\{Y_n\}_n$, denoted by $\{\nu_n\}_n$, $\{\hat{\nu}_n\}_n$ and $\{y_n\}_n$, respectively. We have observed in the previous section that $\mathbf{E}[N_n] = \rho$. On the other hand

$$\mathbf{E}[Y_n] = \mathbf{E}[\mathbf{E}[Y_n | N_n]] = \mathbf{E}[p N_n] = p\rho. \quad (5)$$

Therefore $\nu_n = N_n - \rho$, $\hat{\nu}_n = \hat{N}_n - \rho$ and $y_n = Y_n - p\rho$.

Throughout the paper z is a complex number such that $|z| = 1$. Introduce

$$S_y(z) = \sum_{k=-\infty}^{\infty} \text{Cov}_y(k) z^{-k}$$

the z -transform of the autocorrelation function (also called the power spectrum) of $\{y_n\}_n$.

Let $\text{Cov}_{\nu y}(k) = \mathbf{E}[\nu_{n-k} y_n]$ be the cross-autocorrelation function of processes $\{\nu_n\}_n$ and $\{y_n\}_n$. We also introduce²

$$S_{\nu y}(z) = \sum_{k=-\infty}^{\infty} \text{Cov}_{\nu y}(k) z^{-k}$$

the z -transform of $\text{Cov}_{\nu y}(k)$. We can express $\text{Cov}_{\nu y}(k)$ and $\text{Cov}_y(k)$ in terms of $\text{Cov}_\nu(k)$ as follows

$$\text{Cov}_{\nu y}(k) = p \text{Cov}_\nu(k) \quad (6)$$

$$\text{Cov}_y(k) = \begin{cases} p^2 \text{Cov}_\nu(k), & \text{for } k \neq 0 \\ p^2 \text{Cov}_\nu(0) + \rho p(1-p) = \rho p, & \text{for } k = 0 \end{cases} \quad (7)$$

where we have used the identity $\text{Cov}_\nu(k) = \text{Cov}_N(k)$. We are now in position to derive the Wiener filter. First, we write $S_y(z)$ as

$$S_y(z) = \sigma G(z) G(z^{-1}) \quad (8)$$

where σ is a constant. This operation is called the canonical factorization of the power spectrum of $\{y_n\}_n$. The function $G(z)$ is the part of $S_y(z)$ that has all its zeros and poles inside the unit circle. The function $1/G(z)$ is the transfer function of the linear filter: it transforms $\{y_n\}_n$ into a white noise process with variance σ .

Next, we form the ratio $S_{\nu y}(z)/G(z^{-1})$. This ratio is interpreted as the transfer function of a linear filter. The impulse response of this filter has values at the left and the right of the time origin (non-causal filter). We look for the transfer function of the part of the impulse response at the right of the time origin. This can simply be done by expanding the transfer function into fractions and then taking only the fractions with zeros and poles in the unit circle. In other words, we transfer the filter from a non-causal filter into a causal one. We denote the transfer function of the causal version of the filter by

$$H(z) = \left[\frac{S_{\nu y}(z)}{G(z^{-1})} \right]_+.$$

The transfer function of the optimal filter is then given by

$$H_0(z) = \frac{H(z)}{\sigma G(z)}.$$

²Observe from (6) and (7) that both $S_y(z)$ and $S_{\nu y}(z)$ are well-defined for $|z| = 1$ under the assumption (4).

It remains to invert this transfer function back into the time domain to find the desired recurrence between $\hat{\nu}_n$ and y_n and, subsequently, between the non-centered variables \hat{N}_n and Y_n . This procedure is illustrated in Section 4 in case the underlying model is the $M/M/\infty$ queue.

4 Application of the Wiener filter to the $M/M/\infty$ model

In light of the results reported in Section 3, all what we have to do is to find expressions for $S_y(z)$ and $S_{\nu y}(z)$. This can easily be done when the underlying model is the $M/M/\infty$ queueing model, as shown below.

Let us first determine $S_y(z)$. By using (7) and (3) together with the property that $\text{Cov}_N(k) = \text{Cov}_\nu(k)$, we find

$$\text{Cov}_y(k) = \begin{cases} p^2 \rho \gamma^{|k|}, & \text{for } k \neq 0 \\ p\rho, & \text{for } k = 0. \end{cases}$$

Since $\gamma < 1$ and $|z| = 1$, the z -transform of $\text{Cov}_y(k)$ is

$$\begin{aligned} S_y(z) &= \sum_{k=-\infty}^{-1} p^2 \rho \gamma^{-k} z^{-k} + p\rho + \sum_{k=1}^{\infty} p^2 \rho \gamma^k z^{-k} \\ &= p\rho \left[\frac{\gamma(p-1)z^2 + [1 + \gamma^2(1-2p)]z + \gamma(p-1)}{z(1-\gamma z)(1-\gamma z^{-1})} \right]. \end{aligned}$$

The second-order polynomial in the variable z in the numerator has two positive real roots given by r and $1/r$, with

$$r = \frac{1 + \gamma^2(1-2p) - \sqrt{(1-\gamma^2)[1 - \gamma^2(1-2p)^2]}}{2\gamma(1-p)}.$$

Note that $r < 1$. Hence

$$\begin{aligned} S_y(z) &= \frac{\gamma\rho p(1-p)}{r} \left[\frac{(1-rz)(1-rz^{-1})}{(1-\gamma z)(1-\gamma z^{-1})} \right] \\ &= \sigma G(z) G(z^{-1}) \end{aligned}$$

with

$$\sigma := \frac{\gamma\rho p(1-p)}{r} \quad \text{and} \quad G(z) := \frac{1-rz^{-1}}{1-\gamma z^{-1}}.$$

We now compute $S_{\nu y}(z)$. From (6) and (3) we find

$$\text{Cov}_{\nu y}(k) = p\rho \gamma^{|k|}$$

so that

$$S_{\nu y}(z) = \frac{p\rho(1 - \gamma^2)}{(1 - \gamma z)(1 - \gamma z^{-1})}.$$

The transfer function $H(z)$ is given by

$$H(z) = \left[\frac{S_{\nu y}(z)}{G(z^{-1})} \right]_+ = \frac{\Gamma}{(1 - \gamma r)(1 - \gamma z^{-1})}$$

with $\Gamma := p\rho(1 - \gamma^2)$, and the transfer function $H_0(z)$ of the optimal filter takes here the simple form

$$H_0(z) = \frac{\Gamma}{\sigma(1 - \gamma r)(1 - rz^{-1})}.$$

The impulse response of this linear filter is given by the first-order recurrence relation [8]

$$\hat{\nu}_n = A\hat{\nu}_{n-1} + By_n$$

with $\hat{\nu}_n$ the estimator of ν_n , and

$$A = r, \quad B = \frac{\Gamma}{\sigma(1 - \gamma r)}.$$

We now return to the original processes $\{N_n\}_n$ and $\{Y_n\}_n$, to finally obtain

$$\hat{N}_n = A\hat{N}_{n-1} + BY_n + \rho(1 - A - pB) \tag{9}$$

the best linear filter.

It is interesting to compare this filter with the Kalman filter derived in [3]³. They appear to be the same! This result is not so surprising after all, since both the Kalman filter and the Wiener filter are optimal (among the class of linear filters) in the sense that they minimize the mean square error. The key point is that the Kalman filter used in [3] was derived under a heavy traffic assumption, while the Wiener filter computed in the present paper holds for any value of the model parameters λ and μ . This partly explains why the estimator in [3] behaves well under light or moderate traffic as experimentally observed in that paper.

We conclude this section by computing the mean square error $\epsilon_{min} := \mathbf{E}[(N_n - \hat{N}_n)^2]$ of our estimator. It is known that [8]

$$\epsilon_{min} = \sum_{k=1}^M \text{Res} [F(z), z_k]$$

³Recall that a Kalman filter is the optimal filter under the condition of linear dynamics and observation, which does not hold in our case. However, the dynamics does converge to a linear diffusion as the traffic load tends to infinity, which allowed the authors of [3] to obtain a Kalman filter which is optimal for the asymptotic heavy traffic regime.

with

$$F(z) := \frac{1}{z} [S_\nu(z) - H_0(z)S_{\nu y}(z^{-1})]$$

where z_1, \dots, z_M are the poles (if any) of the function $F(z)$ inside the unit circle. The notation $\text{Res}[F(z), z_k]$ stands for the residue of $F(z)$ at point $z = z_k$, namely, the coefficient of $1/(z - z_k)$ in the Laurent series expansion of $F(z)$ at the vicinity of z_k .

Specializing $F(z)$ to the values of $S_\nu(z)$, $S_{\nu y}(z)$ and $H_0(z)$ found earlier, yields

$$F(z) = \frac{\rho(1 - \gamma^2)}{(1 - \gamma z)(z - \gamma)} \frac{(1 - Bp)z - A}{z - A}.$$

This function has two poles inside the unit circle which are located at $z = A$ and $z = \gamma$; the residues of $F(z)$ at these poles are given by $-\rho pAB(1 - \gamma^2)/((1 - \gamma A)(A - \gamma))$ and $\rho(1 + pB\gamma/(A - \gamma))$, respectively. Summing up these residues gives

$$\epsilon_{min} = \rho - \rho \frac{Bp}{1 - \gamma A}.$$

By using the expressions of A and B we finally obtain

$$\epsilon_{min} = \rho \frac{-(1 - \gamma^2) + \sqrt{(1 - \gamma^2)(1 - \gamma^2(1 - 2p)^2)}}{2\gamma^2 p}. \quad (10)$$

This expression for ϵ_{min} can be used to tune the parameters p and γ or equivalently S (see Section 6).

5 An optimal first-order linear filter

The theory reported in Section 3 applies to any on-time distribution $\Psi(x)$ (with the exception of heavy-tailed distributions). However, it is not easy to identify the function $G(z)$ that appears in the canonical factorization of the spectrum $S_y(z)$ (see (8)) and thereby the optimal filter, except when the join times are exponentially distributed rvs (see Section 4).

In this section we will determine the first-order linear filter that minimizes the mean square error. Observe that, unlike the Wiener filter, the proposed approach will not return the optimal filter among all linear filters but simply the optimal linear filter among all first-order linear filters. We will illustrate this approach at the end of this section in the case where $\Psi(x)$ is an hyperexponential distribution.

Recall the definition of the centered processes $\{\nu_n\}_n$, $\{\hat{\nu}_n\}_n$ and $\{y_n\}_n$ (see beginning of Section 3).

The methodology is simple: we want to find constants $A \in (0, 1)$ and B such that $\epsilon := \mathbf{E}[(\nu_n - \hat{\nu}_n)^2]$ is minimized when the process $\{\hat{\nu}_n\}_n$ satisfies the first-order recurrence relation

$$\hat{\nu}_n = A\hat{\nu}_{n-1} + By_n. \quad (11)$$

In steady-state this implies that

$$\hat{\nu}_n = B \sum_{k=0}^{\infty} A^k y_{n-k}. \quad (12)$$

The mean square error ϵ is equal to

$$\epsilon = \mathbf{E} [\hat{\nu}_n^2] + \mathbf{E} [\nu_n^2] - 2\mathbf{E} [\hat{\nu}_n \nu_n].$$

From (12) and (6) we find

$$\mathbf{E} [\hat{\nu}_n \nu_n] = pB \sum_{k=0}^{\infty} A^k \text{Cov}_\nu(k) = pBg(A)$$

where

$$g(z) := \sum_{k=0}^{\infty} z^k \text{Cov}_\nu(k). \quad (13)$$

The power series $g(z)$ converges for $|z| < 1$ (Hint: $k \rightarrow \text{Cov}_\nu(k)$ is nonincreasing) and is therefore differentiable for $|z| < 1$. We will denote by $g'(z)$ its derivative.

We have $\mathbf{E}[\nu_n^2] = \mathbf{E}[(N_n - \rho)^2] = \rho$ (see Section 2).

It remains to express $\mathbf{E}[\hat{\nu}_n^2]$ in terms of the parameters A and B . Squaring both sides of (11) and then taking the expectation, yields

$$\mathbf{E} [\hat{\nu}_n^2] = \left(\frac{B}{1 - A^2} \right) (2A\mathbf{E} [\hat{\nu}_{n-1} y_n] + B\mathbf{E} [y_n^2]).$$

With the identities $\mathbf{E}[y_n^2] = \text{Cov}_y(0) = \rho p$ (cf. (7)) and $\mathbf{E}[\hat{\nu}_{n-1} y_n] = Bp^2 (g(A) - \rho)/A$ (Hint: use (12), (6) and $\text{Cov}_\nu(0) = \rho$), we obtain

$$\mathbf{E} [\hat{\nu}_n^2] = \left(\frac{pB^2}{1 - A^2} \right) (2pg(A) + \rho(1 - 2p)).$$

Finally, the mean square error is

$$\epsilon = \rho - 2pBg(A) + \left(\frac{pB^2}{1 - A^2} \right) (2pg(A) + \rho(1 - 2p)). \quad (14)$$

In order to minimize ϵ , $A \in (0, 1)$ and B must be the solution of the following system of equations:

$$\begin{cases} \frac{\partial \epsilon}{\partial A} = \frac{2Bp}{1 - A^2} \left(AB \left[\frac{2pg(A) + \rho(1 - 2p)}{1 - A^2} \right] + g'(A)(Bp - (1 - A^2)) \right) = 0 \\ \frac{\partial \epsilon}{\partial B} = \frac{2Bp}{1 - A^2} (2pg(A) + \rho(1 - 2p)) - 2pg(A) = 0. \end{cases}$$

The 2nd equation gives

$$B = \frac{g(A)(1 - A^2)}{2pg(A) + \rho(1 - 2p)}. \quad (15)$$

Reporting this value of B into the 1st equation shows that A must satisfy

$$Ag(A)(2pg(A) + \rho(1 - 2p)) - g'(A)(1 - A^2)(pg(A) + \rho(1 - 2p)) = 0 \quad (16)$$

If this equation has a unique solution $A \in (0, 1)$, then plugging this value of A into (15) will give the optimal pair (A, B) .

It is shown in Appendix A that (16) has always a unique solution in $[0, 1)$ (in particular) if $g'(x) > 0$ in $[0, 1)$. This condition will hold as long as $P(D > S) > 0$. In practice, one can always select S such that this condition is true.

The reader can check that the filter defined in (11) with the optimal pair (A, B) is the same as the Wiener filter found in Section 4 in the case when the on-times are exponentially distributed.

We now illustrate the approach developed in this section by considering the situation where on-times have an hyperexponential distribution. More precisely, we assume that

$$\Psi(x) = 1 - \sum_{i=1}^L p_i e^{-\mu_i x} \quad (17)$$

with $0 < p_i < 1$, $i = 1, 2, \dots, L$, and $\sum_{i=1}^L p_i = 1$. In this setting the underlying queueing model can be seen as L independent $M/M/\infty$ queues in parallel. The arrival rate to queue i is $p_i \lambda$ and the service rate is μ_i . Define $\gamma_i := \exp(-\mu_i S)$, $\rho_i := p_i \lambda / \mu_i$ so that $\rho = \sum_{i=1}^L \rho_i$. The autocorrelation function of the process $\{\nu_n\}_n$ is equal to

$$\text{Cov}_\nu(k) = \begin{cases} \rho & \text{for } k = 0 \\ \sum_{i=1}^L \rho_i \gamma_i^{|k|} & \text{for } k \neq 0. \end{cases}$$

so that

$$g(A) = \sum_{i=1}^L \frac{\rho_i}{1 - A\gamma_i}.$$

*Numerical example*⁴: $L = 2$, $p = 0.0106$ and $S = 2.5s$. Also

$$\begin{array}{lll} 1/\mu_1 & = & 3896.98s \quad \rho_1 = 19.532 \quad \gamma_1 = 0.9993586834 \\ 1/\mu_2 & = & 480061.16s \quad \rho_2 = 75.137 \quad \gamma_2 = 0.9999947923 \\ 1/\mu & = & 18316.32s \quad \rho = 94.669 \end{array}$$

The optimal first-order filter is

$$\hat{N}_n = 0.9987945572 \hat{N}_{n-1} + 0.1072028947 Y_n + 0.00654086406.$$

⁴The values of the parameters come from the trace called *video1* investigated in Section 7. Details on how these values are retrieved come in Appendix B.

For comparison the Wiener filter found in Section 4 (for exponential on-times) for these values is

$$\hat{N}_n = 0.9982858879 \hat{N}_{n-1} + 0.1488534364 Y_n + 0.01290008059.$$

6 Guidelines on how to choose parameters p and S

A “good” pair (p, S) should (i) limit the feedback implosion while in the meantime (ii) achieve a good quality of the estimator. Of course (i) and (ii) are antinomic and therefore a trade-off must be found. This trade-off will be formalized as follows: we want to select a pair (p, S) so that the mean number of ACKs generated every S seconds (see (5)) and the relative error of the variance of the estimator (denoted as η) are bounded from above by given constants, namely

$$\mathbf{E}[Y_n] = p\rho \leq \alpha \quad (18)$$

and

$$\eta = \frac{\text{Var}(N_n) - \text{Var}(\hat{N}_n)}{\text{Var}(N_n)} \leq \beta. \quad (19)$$

When \hat{N}_n is optimal then $\text{Var}(N_n) - \text{Var}(\hat{N}_n) = \mathbf{E}[(N_n - \hat{N}_n)^2]$ and η becomes the “normalized mean square error” [9, p. 202]. Optimality was shown for the $M/M/\infty$ queue and in this case

$$\eta = \frac{\epsilon_{min}}{\rho}$$

where ϵ_{min} is given in (10).

For given constants α and β , it is easy to solve the constrained optimization problem defined in (18) and (19), provided that η is known.

For the $M/M/\infty$ model, where ϵ_{min} is given in (10), we find that $p = \alpha/\rho$ and that S , or equivalently γ , is the unique positive solution of the equation $\epsilon_{min} = \rho\beta$.

The problem now is to choose constants α and β so that conditions (i) and (ii) are satisfied. We have found that α in the range $[0.5, 1]$ and $\beta \leq 0.15$ give satisfactory results.

We conclude this section with general remarks on how to adapt the parameters p and S to important variations in the membership. The estimation schemes in Sections 4 and 5 have been obtained under the assumption that parameters p and S are fixed. However, the filters constructed in Sections 4 and 5 can still be used if p and/or S change over time, provided that these modifications do not prevent the system to be most of the time in steady-state. In that setting, a new filter will have to be recomputed after each modification. Such a modification can be called for each time the number of ACKs received during a given period of time significantly deviates from the current expectation (i.e. $p\rho$).

7 Validating both estimators on real video traces

In this section we apply the estimators developed in Sections 4 and 5 to four real video traces. Two types of estimators will be used: the estimator – denoted as \hat{N}_n^E – found in (9)

when the population is modeled as an $M/M/\infty$ queue; the estimator – denoted as $\hat{N}_n^{H_2}$ – derived in Section 5 in the case when the join times are Poisson and the on-times have a 2-stage hyperexponential distribution ($M/H_2/\infty$ model).

The objective is twofold: we want to investigate the quality of both estimators when compared to real life conditions, and we want to identify the best one.

We have collected four MBone traces – denoted $video_i, i = 1, \dots, 4$ – between August 2001 and September 2001 using the *MListen* tool⁵. Each trace corresponds to a long-lived video session (see duration of each session in Table 1, where the superscript “d” stands for “days”). We have run both algorithms (estimators) on each trace.

For each trace we have identified the parameters of the $M/M/\infty$ model (parameters λ and μ , or equivalently parameters ρ and μ) and of the $M/H_2/\infty$ model (parameters ρ, μ_1, μ_2, p_1 and $p_2 = 1 - p_1$ – see definitions in Section 5). The values of these parameters are reported in columns 3-8 in Table 1.

Parameters p and S have been chosen by following the guidelines presented in Section 6, namely $\alpha \in \{0.5, 1\}$ and $\beta \in \{0.1, 0.15\}$. Values of these parameters are listed in columns 9-10 in Table 1.

	session lifetime	ρ	$1/\mu$	$1/\mu_1$	$1/\mu_2$	p_1	p_2	p	S	α	β
$video_1$	$3^d 13^h 33^m 20^s$	94.7	18316	3897	480061	0.97	0.03	0.011	2.5	1.0	0.15
$video_2$	$11^d 1^h 46^m 8^s$	14.1	16476	1	226498	0.93	0.07	0.034	3.2	0.5	0.1
$video_3$	$50^d 22^h 13^m 20^s$	8.1	66823	1	900854	0.93	0.07	0.062	20.0	0.5	0.1
$video_4$	$29^d 16^h 43^m 13^s$	17.9	83390	1	473268	0.82	0.18	0.028	10.0	0.5	0.1

Table 1: Parameter identification

The performance of estimators \hat{N}_n^E and $\hat{N}_n^{H_2}$ are reported in Tables 2 and 3.

Table 2 reports several order statistics (columns 4-8) and the sample mean of the relative error $\frac{|N_n - \hat{N}_n|}{\hat{N}_n}$ (column 3), where \hat{N}_n is either \hat{N}_n^E or $\hat{N}_n^{H_2}$. All results are expressed in percentage. The first observation is that both estimators perform reasonably well. The sample mean of the relative error is always less than 6.82% and is as low as 3.79%; over all experiments this sample mean is less than 4.5% for both \hat{N}_n^E and $\hat{N}_n^{H_2}$ (see last two lines). The last column gives the 95th percentile and reads as follows: the relative error achieved on trace $video_3$ by \hat{N}_n^E (resp. $\hat{N}_n^{H_2}$) is 95% of the time less than 11.00% (resp. 12.56%). The second observation is that no scheme is uniformly better than the other over an entire session but their sample means are very close to each other (see column 3). For instance, \hat{N}_n^E performs better than $\hat{N}_n^{H_2}$ regarding the 90th and the 95th percentiles whereas the result is reversed regarding the 25th percentile. It looks like the relative error on $\hat{N}_n^{H_2}$ is empirically more dispersed around its mean than is the relative error on \hat{N}_n^E , and it has a longer tail. Across all sessions (see last two lines) 75% of the time $\hat{N}_n^{H_2}$ performs better

⁵ Available at <http://www.cc.gatech.edu/computing/Telecomm/mbone/>

		Mean	25%	50%	75%	90%	95%
<i>video</i> ₁	\hat{N}_n^E	6.82	1.09	2.42	5.25	11.50	19.44
	$\hat{N}_n^{H_2}$	6.12	1.08	2.55	6.31	13.53	20.55
<i>video</i> ₂	\hat{N}_n^E	4.19	1.41	3.08	5.43	8.66	11.91
	$\hat{N}_n^{H_2}$	4.12	0.98	2.14	4.41	8.78	12.55
<i>video</i> ₃	\hat{N}_n^E	4.20	1.55	3.26	5.71	8.71	11.00
	$\hat{N}_n^{H_2}$	3.98	1.07	2.36	4.83	9.35	12.56
<i>video</i> ₄	\hat{N}_n^E	3.79	1.23	2.57	4.51	7.50	10.97
	$\hat{N}_n^{H_2}$	4.06	1.02	2.21	4.39	8.98	14.66
Overall	\hat{N}_n^E	4.44	1.33	2.88	5.22	8.60	11.99
	$\hat{N}_n^{H_2}$	4.34	1.02	2.26	4.73	9.61	14.19

Table 2: Mean and percentiles of the relative error $\frac{|N_n - \hat{N}_n|}{N_n}$

than \hat{N}_n^E . This improvement does not come for free, since it requires the identification of 4 parameters (ρ, μ_1, μ_2, p_1) instead of 2 (ρ and μ) for \hat{N}_n^E .

		Mean	Variance	ϵ_{min}, ϵ	ϵ_{min}/ϵ	η
<i>video</i> ₁	\hat{N}_n^E	-0.1121	12.6641	13.9424		0.1473
	$\hat{N}_n^{H_2}$	-0.0469	12.8508	12.1198	1.150382	
<i>video</i> ₂	\hat{N}_n^E	0.0062	0.4947	1.4068		0.0995
	$\hat{N}_n^{H_2}$	0.0188	0.7851	0.3955	3.557016	
<i>video</i> ₃	\hat{N}_n^E	0.0373	0.2065	0.7370		0.0908
	$\hat{N}_n^{H_2}$	0.0194	0.2291	0.2084	3.536468	
<i>video</i> ₄	\hat{N}_n^E	0.0523	0.9105	1.5656		0.0873
	$\hat{N}_n^{H_2}$	0.0651	1.4231	0.6755	2.317691	

Table 3: Empirical mean and variance of the error $N_n - \hat{N}_n$

Table 3 reports the sample mean and the sample variance of the error $N_n - \hat{N}_n$. Here also \hat{N}_n indifferently denotes \hat{N}_n^E or $\hat{N}_n^{H_2}$. In the 5th column, we have put the theoretical variance. It is given by ϵ_{min} for \hat{N}_n^E (see (10)) and by ϵ for $\hat{N}_n^{H_2}$ (see (14)). The expected average $\mathbf{E}[N_n - \hat{N}_n]$ is zero in both approaches. Both estimators \hat{N}_n^E and $\hat{N}_n^{H_2}$ have almost no bias, and their empirical variances closely match the theoretical ones given by ϵ_{min} and ϵ , respectively. It is of interest to point out that for the 4 traces studied ϵ , the theoretical mean square error provided by $\hat{N}_n^{H_2}$, is smaller than ϵ_{min} , the theoretical mean square error provided by \hat{N}_n^E (however, this result is reversed if we consider the empirical mean square

errors). Thus $\hat{N}_n^{H_2}$ is more efficient⁶ than \hat{N}_n^E (again, \hat{N}_n^E is empirically more efficient than $\hat{N}_n^{H_2}$). The relative efficiency of $\hat{N}_n^{H_2}$ compared to \hat{N}_n^E defined as ϵ_{min}/ϵ is given in the 6th column of Table 3. The last column provides the relative error on $Var(\hat{N}_n^E)$, called η ($= \epsilon_{min}/\rho$) in Section 6.

In Figure 1 (resp. 2, 3 and 4) we have plotted the variations of membership for session *video*₁ (resp. *video*₂, *video*₃ and *video*₄), together with the estimations returned by \hat{N}_n^E and $\hat{N}_n^{H_2}$. Among all 4 sessions, session *video*₁ presented the highest variations in N_n .




Figure 1: Membership estimation of session *video*₁ and probability plots of inter-arrivals and on-times sequences.

Figure 1(a) (resp. 2(a), 3(a) and 4(a)) displays three curves: the collected video trace, the estimation returned by \hat{N}_n^E and labeled “Exponential”, and the estimation returned by $\hat{N}_n^{H_2}$ and labeled “Hyperexponential”. It is clearly visible, especially at the left-hand side of graph 1(a), that \hat{N}_n^E tracks better the session dynamics than $\hat{N}_n^{H_2}$.

⁶An estimator is said to be more efficient if it has a smaller variance.




Figure 2: Membership estimation of session *video₂* and probability plots of inter-arrivals and on-times sequences.

Both estimators \hat{N}_n^E and $\hat{N}_n^{H_2}$ have been derived under some specific and somehow restrictive assumptions: Poisson join times for both of them, exponential (resp. 2-stage hyperexponential) on-times for the first (resp. second) one. It is interesting to know whether or not these assumptions were violated in each session *video_i*, $i = 1, \dots, 4$. We have therefore carried out a statistical analysis of each trace in order to determine the nature of their join time process and of the on-time sequence.

In [1] and [2], Almeroth and Ammar have conducted a similar analysis of various MBone sessions. They have observed that the interarrival (or join) time sequence was well represented by an exponential distribution. As to the on-time sequence, they have noticed two different types of behavior: traces containing long lifetime sessions where the on-time distribution was close to a Zipf distribution, and traces with short lifetime sessions well represented by an exponential distribution.




Figure 3: Membership estimation of session *video*₃ and probability plots of inter-arrivals and on-times sequences.

In contrast with these results, all traces that we have used exhibit a very different behavior, as shown in Table 4 and Figures 1, 2, 3 and 4, parts (b) and (c). For none of them the join time process is Poisson and for none of them the on-times are exponentially distributed (or hyperexponentially distributed). The inter-join times and the on-times appear to follow subexponential distributions (Weibull and lognormal distributions), a situation quite different from the assumptions under which the estimators have been obtained. Despite these significant differences, the estimators behave well and therefore show a good robustness to assumption violations.

In summary, both estimators perform very well when applied to real video traces and are robust to significant deviations from their (theoretical) domain of validity. Estimator $\hat{N}_n^{H_2}$ returns the best global performance for the relative error criterion, but is not as good as \hat{N}_n^E to track high fluctuations. Overall, we have found that \hat{N}_n^E is a good estimator, both in terms of its performance and of the ease with which it can be used since it only requires




Figure 4: Membership estimation of session *video₄* and probability plots of inter-arrivals and on-times sequences.

	Best fit for inter-arrivals sequence	Best fit for on-times sequence
<i>video₁</i>	Lognormal with $\mu = 3.38$, $d = 1.49$	Weibull with shape 0.35, scale 3700
<i>video₂</i>	Lognormal with $\mu = 5.20$, $d = 1.68$	Weibull with shape 0.26, scale 1400
<i>video₃</i>	Weibull with shape 0.65, scale 3500	Lognormal with $\mu = 5.08$, $d = 3.32$
<i>video₄</i>	Weibull with shape 0.55, scale 2700	Weibull with shape 0.18, scale 4000

Table 4: Distributions that best fitted into the inter-arrivals and on-times sequences.

the knowledge of two parameters (λ and μ) that are likely to be known (e.g. from ISPs) in practice.

8 Conclusion

The major contribution of this work is the design of two novel estimators for evaluating the membership in multicast sessions. In contrast to the estimator proposed in [3] which was designed under heavy traffic assumptions, our schemes do not place such restrictions.

Relying on the appealing Wiener filtering theory, we have computed the optimal linear estimator for session membership when the underlying model is an $M/M/\infty$ queue. The optimality refers to the unbiasedness of the estimator and to the fact that the mean square error is minimized. We have also developed the optimal first-order linear filter in case the on-times distribution is arbitrary and have derived the associated estimator in case the on-times have a 2-stage hyperexponential distribution.

Both obtained estimators have been validated on real video traces. Their performance have been shown to be excellent, one of them showing a good ability to adapt to highly dynamic multicast sessions.

A Existence and uniqueness of the solution

Lemma A.1 *Define*

$$f(x) := (2pg(x) + \rho(1 - 2p))xg(x) - (pg(x) + \rho(1 - 2p))(1 - x^2)g'(x)$$

where g is given in (13).

If $g'(x) > 0$ for $x \in [0, 1)$, then f has a unique zero in $[0, 1)$.

Proof. Write $f(x)$ as $f(x) = f_+(x) - f_-(x)$ where

$$\begin{aligned} f_+(x) &:= [2p(g(x) - \rho) + \rho]xg(x) \\ f_-(x) &:= [p(g(x) - \rho) + \rho(1 - p)](1 - x^2)g'(x). \end{aligned}$$

Under the assumptions of the lemma, it is seen that:

- (i) $f_+(x)$ is continuous and strictly increasing in $[0, 1)$ with $f_+(0) = 0$;
- (ii) $f_-(x)$ is continuous in $[0, 1)$ and $f_-(1) = 0$. There exists $x_0 \in (0, 1)$ such that $f_-(x)$ is strictly increasing in $[0, x_0)$, strictly decreasing in $(x_0, 1)$ and $f'_-(x_0) = 0$ (Hint: write $f'_-(x)$ as $-\alpha x^2 - \beta x + \alpha$ with $\alpha := p(g'(x))^2 + [p(g(x) - \rho) + \rho(1 - p)]g''(x) > 0$ and $\beta := 2[p(g(x) - \rho) + \rho(1 - p)]g'(x) > 0$ so that $f''_-(x) < 0$).

We deduce from the above that f has a unique zero in $(0, 1)$. \diamond

B Computing parameters from trace.

Each trace records $(T_i, D_i), i \geq 1$. To use the $M/M/\infty$ queue model, we identify $1/\lambda = \mathbf{E}[T_{i+1} - T_i]$ and $1/\mu = \mathbf{E}[D]$, and deduce $\rho = \lambda/\mu$. To use the $M/H_2/\infty$ queue model, λ is

computed as before. Identifying μ_1, μ_2, p_1 and p_2 requires the knowledge of the first three moments of D (recall that $p_2 = 1 - p_1$). For a 2-stage hyperexponential distribution, the n -th moment is given by

$$\mathbf{E}[D^n] = \sum_{i=1}^2 \frac{p_i n!}{\mu_i^n} = n! \left(\frac{p_1}{\mu_1^n} + \frac{p_2}{\mu_2^n} \right), \text{ for } n \geq 1.$$

It then follows that $1/\mu_1$ and $1/\mu_2$ are the (positive) solutions of

$$6\left(2\mathbf{E}[D]^2 - \mathbf{E}[D^2]\right)x^2 + 2\left(\mathbf{E}[D^3] - 3\mathbf{E}[D^2]\mathbf{E}[D]\right)x + 3\mathbf{E}[D^2]^2 - 2\mathbf{E}[D^3]\mathbf{E}[D] = 0.$$

Finally,

$$\begin{aligned} p_1 &= \frac{\mathbf{E}[D]/\mu_2 - \mathbf{E}[D^2]/2}{1/\mu_1(1/\mu_2 - 1/\mu_1)} \\ p_2 &= \frac{\mathbf{E}[D^2]/2 - \mathbf{E}[D]/\mu_1}{1/\mu_2(1/\mu_2 - 1/\mu_1)} = 1 - p_1. \end{aligned}$$

We now can compute $\rho_i = p_i \lambda / \mu_i$ and $\gamma_i = \exp(-\mu_i S)$ for $i = 1, 2$. At last $\rho = \rho_1 + \rho_2$.

References

- [1] 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*, pages 209–216, August 1996.
- [2] K. Almeroth and M. Ammar. Multicast group behavior in the Internet's Multicast Backbone (MBone). *IEEE Communications Magazine*, 35:224–229, June 1997.
- [3] S. Alouf, E. Altman, and P. Nain. Optimal on-line estimation of the size of a dynamic multicast group. *To appear in Proc. of IEEE INFOCOM'2002, New York, NY USA*.
- [4] J.-C. Bolot, T. Turetti, and I. Wakerman. Scalable feedback control for multicast video distribution in the Internet. In *Proc. of ACM SIGCOMM'94, London, UK*, pages 58–67, September 1994.
- [5] D. R. Cox and V. Isham. *Point Processes*. Chapman and Hall, New York, 1980.
- [6] 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, NY USA*, pages 342–356, 1995.
- [7] T. Friedman and D. Towsley. Multicast session membership size estimation. In *Proc. of IEEE INFOCOM'99, New York, NY USA*, volume 2, pages 965–972, March 1999.

-
- [8] S. Haykin. *Modern Filters*. Macmillan, New York, 1989.
 - [9] S. Haykin. *Adaptive Filter Theory*. Prentice Hall, 3rd edition, 1996.
 - [10] L. Kleinrock. *Queueing Systems: Theory*, volume 1. John Wiley and Sons, 1975.
 - [11] C. Liu and J. Nonnenmacher. Broadcast audience estimation. In *Proc. of IEEE INFOCOM 2000, Tel Aviv, Israel*, volume 2, pages 952–960, March 2000.
 - [12] J. Nonnenmacher and E. Biersack. Optimal multicast feedback. In *Proc. of IEEE INFOCOM'98, San Francisco, CA USA*, volume 3, pages 964–971, March 1998.
 - [13] J. Nonnenmacher and E. Biersack. Scalable feedback for large groups. *IEEE/ACM Trans. on Networking*, 7(3):375–386, June 1999.
 - [14] J. Rosenberg and H. Schulzrinne. Timer reconsideration for enhanced RTP scalability. In *Proc. of IEEE INFOCOM'98, San Francisco, CA USA*, volume 1, pages 233–241, March/April 1998.
 - [15] H. Schulzrinne, S. Casner, R. Frederick, and V. Jacobson. RTP: a transport protocol for real-time applications. RFC 1889, Network Working Group, January 1996.



Unité de recherche INRIA Sophia Antipolis
2004, route des Lucioles - BP 93 - 06902 Sophia Antipolis Cedex (France)
Unité de recherche INRIA Lorraine : LORIA, Technopôle de Nancy-Brabois - Campus scientifique
615, rue du Jardin Botanique - BP 101 - 54602 Villers-lès-Nancy Cedex (France)
Unité de recherche INRIA Rennes : IRISA, Campus universitaire de Beaulieu - 35042 Rennes Cedex (France)
Unité de recherche INRIA Rhône-Alpes : 655, avenue de l'Europe - 38330 Montbonnot-St-Martin (France)
Unité de recherche INRIA Rocquencourt : Domaine de Voluceau - Rocquencourt - BP 105 - 78153 Le Chesnay Cedex (France)

Éditeur
INRIA - Domaine de Voluceau - Rocquencourt, BP 105 - 78153 Le Chesnay Cedex (France)
<http://www.inria.fr>
ISSN 0249-6399