On the dynamic estimation of multicast group sizes

Sara Alouf, Eitan Altman, Chadi Barakat and Philippe Nain

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

Abstract

This paper concerns multicast applications that are interested in the evolution of their membership over time. It covers optimal on-line estimation algorithms for determining the membership of a multicast group. The paper briefly reviews the related work and our own contributions to the field. Using a probabilistic acknowledgement scheme and signal processing's filtering techniques, we have derived MSE-optimal estimators under the assumptions of Poisson subscribers arrival and either exponentially or hyperexponentially distributed lifetime of receivers. Our estimators have been tested through trace-driven simulations using data from real multicast video sessions over which they exhibit very good performance.
Keywords: on-line estimation, multicast, $M/M/\infty$ queue, diffusion, Kalman filter, Wiener filter.

1 Introduction

Since its introduction, IP multicast [6] has seen slow deployment in the Internet, as the service model and architecture do not efficiently provide many features required for a robust implementation. However, the fact remains that IP multicast is very appealing in offering scalable point-to-multipoint delivery specially in satellite communications and for large-scale multicast applications such as real-time stock quote dissemination, live sports video feeds or Internet radio and TV.

This work has been motivated by the conviction that large-scale multicast applications will be widely deployed in the future as soon as the capability becomes available. We believe that membership estimates will be an essential component of this widespread deployment as they can be very useful for scalable multicast. For instance, the membership of a session can be used for feedback suppression or for charging the sources in large-scale applications. ISPs traditionally charge their customers on an input-rate basis. An alternative pricing scheme would be to charge sources based on their audience size which is more profitable in the case of millions of subscribers. Also, estimating the size of a multicast session can be quite useful to many applications. As an example, Bolot *et al.* [5] use membership estimation to further estimate the proportion of congested receivers as needed for their videoconference system IVS [10].

There has been a significant research effort in devising sampling-based schemes for the estimation of the membership in multicast sessions [5, 7, 13, 14]. The feedback algorithms presented in these references are all *at-least-one* scenarios in the sense that the membership estimation is based on at least one acknowledgement (ACK) coming from the receivers. In these probabilistic schemes, the receivers send ACKs to the source as a reply to a specific request, either with a certain probability as in [5], or after some random time like in [7, 13, 14]. But what is common to these schemes – except the one used in [14] – is that they all assume that the size of the group does not change during the estimation process. Whenever an estimation of the population size is needed, the application re-runs the estimation algorithm without taking into account previous estimates.

Our contribution to the field relies on a novel sampling-based technique which is not an *at-least-one* scenario. Whenever a source is interested in knowing how many recipients are connected to the multicast session, it asks all of the connected members to send an ACK with some *small* probability p[1]. Yet, p needs to be chosen carefully because if it is too small, the estimation would be inaccurate. To be able to track the population size, the source should ask the receivers to repeatedly send ACKs say every S seconds. Occasionally, the source re-issues this request to insure that newly arrived receivers participate in the polling. If S is not too large then the population size at two consecutive estimation instants would present statistical dependence. In order to benefit from this dependence the best possible way, for instance to obtain a better estimation or to get a given quality of estimation with a smaller required volume of ACKs (i.e. decreasing p or increasing S), we use signal processing tools such as the powerful Kalman filter or the Wiener filter to process the ACKs collected at the source.

In the following, we will overview *(i)* in brief, the sampling-based estimation schemes studied in the literature and *(ii)* our approach based on adaptive filter theory.

[1] It is undesirable to have all receivers reply, specially in case of large populations, as the ACKs could overload the network. This undesirable event is called *feedback implosion*.

2 Related work

In the previous section, we briefly introduced the different feedback mechanisms proposed in the literature. We will now review each technique separately, following their chronological conception.

The feedback mechanism proposed by Bolot *et al.* in [5] (called **BTW** mechanism) consists of a series of probabilistic polling rounds, each with a higher reply probability than in the previous one, until feedback is obtained. Each round begins when the sender (or source) multicasts a polling request in which the reply probability p_n is specified. In the first round, $p_1 = 2^{-16}$, and for each subsequent round n, $n \geq 1$, we have $p_n = 2^{n-2}/(2^{16} - 2^{n-2})$. After issuing a polling request, the sender sets a timer at twice the largest round trip time in the receiving group. When the timer times out, the sender initiates a new round with a higher probability. The polling rounds keep going until either a reply has been received, or the round in which the reply probability is 1 has been reached, in which case any receiver will send a response. This ends the ongoing series of rounds. The authors of [5] map experimentally the number of receivers N to the average round $\mathbf{E}[\text{First}]$ in which the first reply is received as follows:

$$N \sim e^{16.25 - \mathbf{E}[\text{First}]/1.4} = \hat{N}_{BTW}. \tag{1}$$

Even though, this mechanism is scalable and avoids feedback implosion (unless N is an order of magnitude greater than 2^{16}), the estimator deriving from (1) does not fully render the variations in the membership.

In [14, 15], Nonnenmacher and Biersack deeply analyze timer-based schemes for multicast feedback. They evaluate the performance of this family of mechanisms, their main concern being the scalability to groups as large as 10^6 receivers. The feedback implosion problem is handled at the receivers: each participant multicasts his response unless he receives one from another participant, in which case he will suppress his own feedback. The timer-based feedback proposed by Nonnenmacher and Biersack (called **NB** mechanism) is round-based and works as follows. Based on an estimation of N, the number of receivers, the sender computes λ and T, the parameters of a truncated exponential distribution (each timer coming from this distribution is in the interval $[0, T]$). At the beginning of each round n, the source multicasts a request for feedback (n, λ, T). The receivers set their timers accordingly and send a feedback message when the timer times out, unless it is suppressed by another message. On the receipt of the feedback messages, the sender estimates N using the timer settings of all of the receivers that returned feedback, which triggers the computation of λ and T for the subsequent round. To express the estimator of N, let $F(z)$ be the distribution of the truncated exponential timer z, c be the *constant* delay between receivers and between any receiver and the sender, m be the minimal timer among the feedback returned and Y be the amount of feedback returned. It is therefore shown [14] that

$$F(z) = \frac{e^{\lambda z/T} - 1}{e^{\lambda} - 1},$$
$$\hat{N}_n = \frac{Y[1 - F(m)]}{F(m+c) - F(m)} = Y\frac{e^{\lambda(1-m/T)} - 1}{e^{\lambda c/T} - 1},$$
$$\hat{N}_{n,\alpha} = \begin{cases} 1, & n = 1, \\ \alpha \hat{N}_{n-1,\alpha} + (1-\alpha)\hat{N}_n, & n > 1. \end{cases} \tag{2}$$

The exponential weighted moving average expression $\hat{N}_{n,\alpha}$, with $\alpha = 0.8$, is used as membership estimator, whereas the expression of $\hat{N}_n$ is used to compute λ and T to react faster to changes in the population size. It readily comes from (2) that $\mathbf{E}[\hat{N}_{n,\alpha}] = \mathbf{E}[\hat{N}_n]$ in steady-state, but it is seen that $\mathbf{E}[\hat{N}_n] \neq N$, and the bias on the estimator depends on the membership N as illustrated in [7]. Another issue is the choice of the parameter α in (2). Nonnenmacher and Biersack suggest the use of $\alpha = 0.8$ to achieve a fast convergence and a reasonably smooth estimate. However, it is not known whether this choice is optimal or not.

To the best of our knowledge, Friedman and Towsley were the first ones to investigate the estimation of the membership size as a whole. In [7], they base their analysis upon a mapping of the polling mechanisms 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 [5, 14] which have point estimators and further add some contributions to each. The reproductions of their main results is beyond the scope of this paper, the interested reader is therefore strongly advised to refer to [7] for full informations or to [1, page 68] for a summary.

Another timer-based feedback scheme is proposed in [13] in which receivers send their randomly delayed reply 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 versions, there is a risk of a feedback implosion. In each version, a maximum likelihood estimator is derived and multiple polling rounds are necessary to return a single estimate. The paper focused on the quality of the estimator rather than on the dynamic nature of multicast sessions.

To conclude this section on related works, we would like to briefly discuss a paper "On the scaling of feedback algorithms for very large multicast groups" [8]. This paper does not deal with the estimation of the membership itself, but rather on its impact on three feedback algorithms, which all are *at-least-one* scenarios. It is concluded that the possible estimations of the group size might be a source for disturbances and that the *exponential feedback raise* algorithm is the algorithm of choice for very large groups.

Both the **BTW** and **NB** mechanisms are considered to have an exponential feedback raise.

3 Motivation

In order to fully reproduce the evolution of the multicast membership, we aim at developing a moving average estimator like the one in (2). That estimator was shown to be biased [7], while what we are actually looking for is an *unbiased* estimator that would take advantage of previous estimates in an *optimal* way.

In [4], we have proposed a mechanism in which the receivers probabilistically send "heartbeats" to the sender in a periodic way. The feedback implosion problem is addressed via a convenient choice of the reply (or ACK) probability p. The ACK interval S between two consecutive polling instants has to be larger than the largest round-trip time between a receiver and the source. Inter-hosts delays are not required to be homogeneous because the ACK interval S is large enough in order to have all of the ACKs produced in a round reach the source before the (automatic) start of the next round. It is explained in [3] how the interval S and the probability p can be set so as to achieve high quality estimation while simultaneously avoiding feedback implosion.

A simple approach to process the amount of ACKs Y_n received in a round n consists of using an exponential weighted moving average (EWMA) like the one used in (2). We can write

$$\hat{N}_n = \alpha \hat{N}_{n-1} + (1-\alpha) Y_n / p. \tag{3}$$

In steady-state, we have $\mathbf{E}[\hat{N}_n] = \mathbf{E}[Y_n]/p = \mathbf{E}[N_n]$. The estimator $\hat{N}_n$ is then unbiased. But the problem lies in the choice of the parameter α. Highly dynamic sessions would require a relatively low α, whereas tracking slowly varying populations requires that α be close to 1. The correct approach is to compute the optimal α that minimizes the estimation error. Notice that (3) is an autoregressive equation of the form $\hat{N}_n = A\hat{N}_{n-1} + BY_n$. One might wonder if this form is the best one or not. Instead of computing the optimal α for this particular form, we have relied on adaptive filter theory to construct the best estimator.

Remark 3.1 *Throughout this paper, it is assumed that neither the requests for ACKs sent by the source, nor the ACKs sent by the receivers, are lost. The loss of polling requests has a smaller impact on the membership estimation mechanism as the source can repeatedly send them or send a group of them whenever the parameters S and/or p are to be changed. As for the loss of ACKs, it is possible to incorporate the loss probability in our feedback mechanism. Let p_L denote the probability that an ACK is lost before it reaches the source. It suffices then to require an ACK probability of $p/(1-p_L)$ in order to have, over the session duration, an average of $p\mathbf{E}[N]$ ACKs received per round. This is equivalent to requiring an ACK probability of p in a safe environment (i.e. no losses).*

4 Optimal estimation

In [4], we have derived the MSE-optimal estimator using a Kalman filter. We have used some simplifying assumptions that allowed us to obtain a good estimation scheme, which, even if not always the optimal, shows good performance. To that end we have considered an exponential distribution for the time during which a receiver stays in the multicast session (referred to as "lifetime" or even "on-time") and made a large group size assumption. This allowed us to obtain a diffusion approximation for the population dynamics. Sampling this process at some regular time intervals yields a discrete-time linear stochastic difference equation for the population dynamics. We have further derived a linear discrete-time equation for the measurements. The fact that both the population dynamics and the measurements in our approximations are linear, have allowed us to use the Kalman filtering theory to design a simple dynamic estimation procedure which is optimal for the heavy traffic model (in minimizing the second moment of the error). This scheme thus makes the best use of previous estimates in order to update the current estimation optimally.

Letting T_i and $T_i + D_i$ be the join time and the leave time, respectively, of the ith participant to a multicast group, the number of participants (or session size) at time $t \geq 0$, $\tilde{N}(t)$, is simply

$$\tilde{N}(t) = \sum_{i=1}^{\tilde{N}(0)} \mathbf{1}\{D_i^{(r)} > 0\} + \sum_{i=1}^{\infty} \mathbf{1}\{T_i \leq t < T_i + D_i\}$$

where $\{D_i^{(r)}, i = 1, 2, \ldots, \tilde{N}(0)\}$ are the remaining on-times at $t = 0$ of participants, if any, which have joined the session before $t = 0$ and who are still connected at time $t = 0$, and $\mathbf{1}\{E\}$ is the indicator function of the event E. In this paper, we work only with the stationary version of the process $\{\tilde{N}(t), t \geq 0\}$, denoted by $\{N(t), t \geq 0\}$.

The times $t = nS, n = 1, 2, \ldots$, will denote the end of each round, so that at time nS, the source possesses all of the ACKs sent to it by participants in the nth round. Throughout, 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 $n = 1, 2, \ldots$.

Primarily for mathematical tractability, we have assumed in [4] that the arrival process is Poisson with rate $\lambda T > 0$ and that on-times form a renewal sequence with common exponential distribution with finite mean $1/\mu$, further independent of the arrival process, so that the process $\{N(t), t \geq 0\}$ is simply the occupation process in a $M/M/\infty$ queueing system with arrival rate λT and mean service time $1/\mu$ [12]. We have investigated in [4] the $M/M/\infty$ queue in heavy traffic, meaning that $T \to \infty$. Since $N_T(t) \to \infty$ a.s. as $T \to \infty$, we have introduced the normalized process $\{Z_T(t), t \geq 0\}$ defined by

$$Z_T(t) = \frac{N_T(t) - \lambda T/\mu}{\sqrt{T}}, \quad t \geq 0, \tag{4}$$

which converges to the Ornstein-Ühlenbeck process $\{X(t),\, t \geq 0\}$ as $T \to \infty$ (see [16, Theorem 6.14, page 155]

$$X(t) = e^{-\mu t}\, X(0) + \sqrt{2\lambda} \int_0^t e^{-\mu(t-u)}\, \mathrm{d}B(u),$$

where $\{B(t),\, t \geq 0\}$ is the standard Brownian motion. We know from [11, page 358] that $X(t)$ is an ergodic Markov process and its invariant distribution is a normal distribution with mean zero and variance $\rho = \lambda/\mu$. Let $\gamma = \exp(-\mu S)$. It is seen in [4] that $X(t)$ satisfies the following equation

$$X((n+1)S) = \gamma\, X(nS) + w_n, \quad n = 0, 1, \ldots \tag{5}$$

where $w_n := \sqrt{2\lambda} \int_{nS}^{(n+1)S} e^{-\mu((n+1)S-u)}\, \mathrm{d}B(u)$ is a zero mean Gaussian noise with variance $Q = \rho\,(1-\gamma^2)$.

At the other hand, denote by Y_n the number of ACKs received in round n and introduce the normalized process

$$M_T(nS) = \frac{Y_n - p\rho T}{\sqrt{T}}, \quad n = 0, 1, \ldots, \tag{6}$$

which, with the help of (4), can be rewritten as

$$M_T(nS) = p\, Z_T(nS) + \frac{Y_n - N_T(nS)p}{\sqrt{T}}.$$

We have shown in [4] that $M_T(nS)$ converges weakly as $T \to \infty$ to a random variable $m(nS)$ such that

$$m(nS) = pX(nS) + v_n, \quad n = 0, 1, \ldots, \tag{7}$$

where v_n is a zero mean Gaussian noise with variance $R := \rho\, p\,(1-p)$. A detailed proof is available in [1, Ch. 2, page 76].

Equations (5) and (7) represent the equations of a discrete time linear filter, for which we can compute the optimal estimator. Let $\hat{X}_n$ be an estimator of $X(nS)$. The estimator that minimizes the mean square of the estimation error is given by the following Kalman filter (see e.g. [18, page 347]), which, in its stationary version, has the following structure:

$$\begin{aligned} P &= \left(\left(\gamma^2\, P + Q\right)^{-1} + p^2/R \right)^{-1} \\ K &= Pp/R \\ \hat{X}_n &= \gamma \hat{X}_{n-1} + K \left[m(nS) - p\gamma \hat{X}_{n-1} \right] \end{aligned} \tag{8}$$

where the constants R and Q have been defined earlier in the section, P gives the steady-state variance of the estimation error and K is the filter gain. The latter can be expressed in terms of p and $\gamma = \exp(-\mu S)$ as follows:

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

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. Motivated by (4), we define $\hat{N}_n$ as follows:

$$\hat{N}_n = \hat{X}_n \sqrt{T} + \rho T.$$

In other words (use (6) and (8)), we have the following estimate for $N_T(nS)$

$$\hat{N}_n = \gamma(1-Kp)\hat{N}_{n-1} + K\, Y_n + \rho T(1-\gamma)(1-Kp), \tag{9}$$

which is unbiased[2] (see [4]) in the sense that

$$\mathbf{E}\left[\hat{N}_n - N_T(nS) \right] = 0.$$

Alternative approach

In [2], we have derived the same estimator, given in (9), but under slightly more general assumptions. Indeed, we have used Wiener filter theory to derive the optimal estimator. The dynamics are no longer required to be linear as in the case of the Kalman filter, allowing us to remove the heavy traffic assumption made earlier. The intensity of the arrival process is now λ instead of λT. In [2], we have introduced centered version of all processes at hand, namely, the membership process $\{N(nS)\}_n$ and the measurement process $\{Y_n\}_n$ to be able to use the Wiener filter. We have used the *prewhitening approach* [9, page 81] to compute the transfer function of the optimal filter which takes as input the centered measurement process $\{y_n\}_n$ ($y_n := Y_n - p\rho$) and gives as output the estimation $\{\hat{\nu}_n\}_n$ of the centered membership process $\{\nu_n\}_n$ ($\nu_n := N(nS) - \rho$). The membership estimator is then defined as $\hat{N}_n = \hat{\nu}_n + \rho$, which is exactly the one in (9). Observe that the Wiener filter approach can be used when the on-time distribution is general ergodic, but an explicit formulation of the transfer function is possible only for the case when on-times are exponentially distributed.

In a second attempt to generalize the model, we have developed in [2] an efficient estimator under the assumptions of Poisson arrivals and hyperexponentially distributed on-times. We have relied on the 'Least Mean Squares' algorithm to derive the first-order linear filter, which is optimal among the class of all first-order linear filters, unlike the Wiener filter which is the optimal (first-order) linear filter among the class of all linear filters, and unlike the Kalman filter which is, under normality assumptions, optimal among all measurable filters and not only among all linear filters based on a set of observations [17, 18].

All the estimators that we have derived have been tested through trace-driven simulations using both synthetic and real traces that do not verify the assumptions under which the estimators were derived. The reader is referred to [1, Ch. 2] for all simulations results and analysis.

Observe that to be able to use the estimators in practice, one should know, or otherwise estimate, the arrival intensity λ and the expected membership ρ. Using very simple

[2] The estimator would be asymptotically unbiased if we have not considered the stationary process $\{N_T(t), t \geq 0\}$.

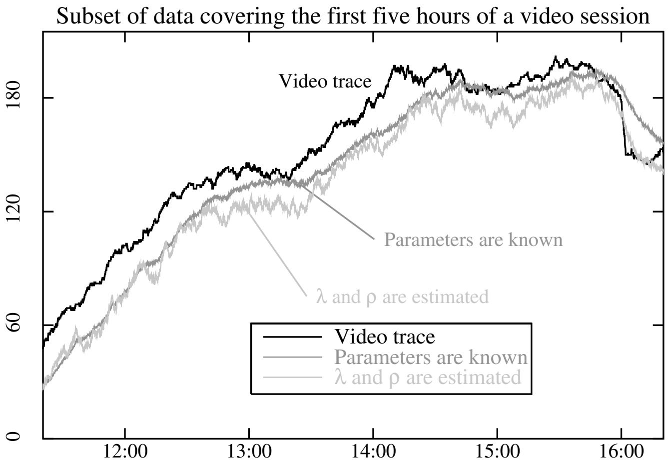

Figure 1: Membership estimation of a video session when parameters are (*i*) known, (*ii*) estimated

estimators for the latter parameters (detailed in [3]), we apply equation (9) to estimate the membership of a multicast video session. The performance of the estimator can visually be observed in Figure 1 in which three curves are plotted: (*i*) the membership in the original video trace, (*ii*) the membership estimation for the case where the parameters are known beforehand, (*iii*) the membership estimation for the case where estimators for λ and ρ are used. We have represented only a subset of the data which corresponds to the first five hours, since the beginning of the session is the most challenging for our algorithm. As expected, when the parameters ρ and λ are unknown, the estimator $\hat{N}_n$ does not behave as well as when the parameters are known beforehand. Still, its performance is reasonably fair as can be seen in Figure 1.

5 Conclusion

This paper reviews the multiple probabilistic techniques found in the literature that enables the estimation of the membership in multicast sessions, whether this estimation is meant to be for a static or dynamic session. After a brief overview of the related work, we have presented our contributions to the field which make use of robust filtering techniques, such as the Kalman filter and the Wiener filter.

References

[1] S. Alouf. *Parameter Estimation and Performance Analysis of Several Networking Applications.* PhD thesis, University of Nice Sophia-Antipolis, November 2002. Available at www.inria.fr/maestro/personnel/Sara.Alouf/these.html.

[2] S. Alouf, E. Altman, C. Barakat, and P. Nain. Estimating membership in a multicast session. *Performance Evaluation Review, (Proc. of ACM SIGMETRICS '03, San Diego, California)*, 31(1):250–260, June 2003.

[3] S. Alouf, E. Altman, C. Barakat, and P. Nain. Optimal estimation of multicast membership. *IEEE Trans. on Signal Processing, Special Issue on Signal Processing in Networking*, 51(8):2165–2176, August 2003.

[4] S. Alouf, E. Altman, and P. Nain. Optimal on-line estimation of the size of a dynamic multicast group. In *Proc. of IEEE INFOCOM '02, New York, New York*, volume 2, pages 1109–1118, June 2002.

[5] 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*, pages 58–67, September 1994.

[6] S. Deering. *Multicast Routing in a Datagram Internetwork.* PhD thesis, Stanford University, December 1991.

[7] T. Friedman and D. Towsley. Multicast session membership size estimation. In *Proc. of IEEE INFOCOM '99, New York, New York*, volume 2, pages 965–972, March 1999.

[8] T. Fuhrmann and J. Widmer. On the scaling of feedback algorithms for very large multicast groups. *Computer Communications, Special Issue on Integrating multicast into the Internet*, 24(5-6):539–547, March 2001.

[9] S. Haykin. *Modern Filters.* Macmillan, New York, 1989.

[10] INRIA. *IVS*, 1992. http://www.inria.fr/rodeo/ivs.html.

[11] I. Karatzas and S. E. Shreve. *Brownian Motion and Stochastic Calculus.* Springer, 1991.

[12] L. Kleinrock. *Queueing Systems: Theory*, volume 1. John Wiley and Sons, 1975.

[13] C. Liu and J. Nonnenmacher. Broadcast audience estimation. In *Proc. of IEEE INFOCOM '00, Tel Aviv, Israel*, volume 2, pages 952–960, March 2000.

[14] J. Nonnenmacher. *Reliable multicast transport to large groups.* PhD thesis, Ecole Polytechnique Federale de Lausanne, Switzerland, July 1998.

[15] J. Nonnenmacher and E. Biersack. Scalable feedback for large groups. *IEEE/ACM Trans. on Networking*, 7(3):375–386, June 1999.

[16] P. Robert. *Réseaux et Files d'Attente: Méthodes Probabilistes*, volume 35 of *Mathématiques & Applications.* Springer Paris, 2000.

[17] P. Ruckdeschel. *Robust Kalman Filtering – Optimality of the Kalman Filter.* MD*TECH Method and Data Technologies, June 2002. http://www.xplore-stat.de/tutorials/rkalmframe3.html.

[18] R. F. Stengel. *Stochastic Optimal Control, Theory and Application.* John Wiley & Sons, 1986.