A *Strike Limit Algorithm* based on Adaptive Kalman Filtering with an application to aboriginal whaling of bowhead whales

Eva Hlín Dereksdóttir and Kjartan G. Magnússon

Science Institute, University of Iceland, Dunhagi 3, 107 Reykjavík, Iceland
Contact e-mail: ehd@ki.is

ABSTRACT

A full and detailed description of a *Strike Limit Algorithm* (*SLA*) based on Adaptive Kalman Filtering techniques with an application to the Bering-Chukchi-Beaufort (B-C-B) Seas stock of bowhead whales is presented in this paper. Extended Kalman filters are used to estimate the present stock size and posterior probability distributions for *MSY*-rate (*MSYR*) and the pre-exploitation stock size *K*. A catch control law selected from a one-parameter family of such rules is then used on the conditional estimates of stock size. These conditional strike limits together with the posterior distributions of the various combinations of *MSYR* and *K*, give a cumulative distribution function for the strike limit. The eventual strike limit is then determined as a pre-specified percentile of this distribution. The *SLA* can be tuned to varying degrees of risk by the choice of the parameter characterising the catch control law and the percentile of the distribution for the strike limit. The procedure is tested on the *Evaluation Trials* set by the Standing Working Group on Aboriginal Whaling Management Procedures.

INTRODUCTION

Non-profit hunting of large whales by native communities where the objective is to meet the nutritional and cultural requirements of the hunting communities is managed by the International Whaling Commission as aboriginal subsistence whaling (e.g. Donovan, 1982). A typical example is the hunting of bowhead whales *(Balaena mysticetus)* by Alaskan Inuit communities (Stoker and Krupnik, 1993). Any type of harvesting requires some type of management and the Commission of the IWC has specified the following management objectives (IWC, 1999):

(1) ensure that the risks of extinction to individual stocks are not seriously increased by subsistence whaling;
(2) enable aboriginal people to harvest whales in perpetuity at levels appropriate to their cultural and nutritional requirements, subject to the other objectives; and
(3) maintain the status of stocks at or above the level giving the highest net recruitment and to ensure that stocks below that level are moved towards it, so far as the environment permits.

In connection with the second objective, the annual number of whales needed to meet the requirements of the hunting communities must be specified. It is agreed by the IWC in practice for each harvest. This number will hereafter be referred to as aboriginal need. It may be constant or time-varying, but is always specified in advance and forms an upper bound on the number of whales which can be taken. The management objectives may not be compatible since fulfilling need may entail an increased risk of extinction. A trade-off between objectives is often inevitable, and good management is thus a compromise between the different objectives. In order to manage a stock successfully, a fully specified management procedure or management scheme has many advantages (e.g. see discussions of the IWC's Revised Management Procedure for commercial whaling of baleen whales: Hammond and Donovan, In press). Data requirements, guidelines for data collection and data treatment, rules whereby catch limits are set based on input data, rules for dealing with the absence of data, time period for which catch limits are set, and so on, are all fully specified and defined in such a scheme. The *ad hoc* nature of assessments and the setting of quotas or strike limits is therefore removed and the whole process becomes more or less automated. A central component of a management scheme is a *Strike Limit Algorithm* (*SLA*), which is a rule or an algorithm where a data series − usually abundance data − is input into the algorithm, and the output is the number of whales which may be taken or struck. The performance of an *SLA* may be tested and evaluated prior to application to real stocks by computer simulations using mathematical models of stock dynamics, abundance observations and data quality. The robustness of the *SLA* can be tested for a wide variety of scenarios by altering the assumptions about parameter values and functional forms in the simulation model. A degree of confidence that the *SLA* can safely be applied to real stocks can therefore be obtained in advance. *SLA*s can take many different forms, but usually contain an estimation part and a part for setting strike limits. In this paper an *SLA* based on the Adaptive Kalman Filter (AKF – see Dereksdóttir and Magnússon, 2001) is presented and an application to aboriginal whaling is given. This *SLA* will be referred to as the AKF-*SLA*.

Kalman filters (Kalman, 1960) are widely used to estimate the state of a stochastic dynamical system with noisy observations, i.e. a system with both 'process noise' and 'observation noise'. Kalman filtering has been applied to estimation problems in fisheries management with some success (e.g. Pella, 1993; Gudmundsson, 1994; Kimura *et al.*, 1996; Reed and Simons, 1996) to estimate stock sizes and population parameters using catch and effort data. The equations used to describe the population dynamics of whales and the assumed relation between the true stock size and observations thereof can be written in a form suitable for state estimation by Kalman filters. Such estimation schemes together with a set of catch control laws form the basis of the *SLA* presented here.

In order to apply the Kalman filter, mathematical models of the state dynamics and the relationship between the observations and the true state are required. The way the Kalman filter works is that the most recent state estimate is projected forward in time by the model and the next observation predicted. The prediction is then compared to the actual observation and the state estimate corrected. The correction or update is proportional to the difference between the prediction and the observation. A large difference therefore results in a large correction and a small difference results in a correspondingly small update in the estimate. The proportionality constant, known as the Kalman gain, depends on the magnitude of the measurement noise and the noise in the dynamics. A high value of the observation noise – meaning that the level of confidence in the observations is low – gives a small gain and thus a small correction in the model prediction. On the other hand the gain is high if the measurement noise is small relative to the process noise and the level of confidence in the observations therefore high. The updated estimate of the state is then projected forward in time by the model until a new observation is made and the process repeated. In the Kalman filtering application presented in this paper, the state of the system is the total (1+) population size of the stock, the component on which density-dependence is assumed to operate, and the observations are the survey estimates. For a fuller description of Kalman filtering methods and their properties, see for example Brown and Hwang (1997).

The version of the AKF-*SLA* presented in this paper is adapted to the Bering-Chukchi-Beaufort Seas stock of bowhead whales for which a considerable amount of information is available, including a series of abundance observations. Furthermore, it is expected that observations will continue to be made at regular intervals in the future. A set of simulation trials designed to evaluate the performance of potential *SLA*s for this stock is given in IWC (2003a). These *Evaluation Trials* are conditioned[1] on the data for the B-C-B bowhead stock, i.e. on catch history, past stock estimates and parameter values.

The *SLA* described here forms a part of the Aboriginal Whaling Management Procedure for the B-C-B stock of bowhead whales adopted by the Scientific Committee of the IWC in 2002 (IWC, 2003b). This may in fact be the first example where management will actually be based on a Kalman filter approach. This paper contains a full and detailed description of the algorithm: in the next section a general framework is given so that the estimation methods of Kalman filtering can be applied, followed by a fully specified model applicable to the B-C-B stock. The results of applying this *SLA* to the *Evaluation Trials* are given and finally, the results of some sensitivity tests are presented and discussed.

MODEL FORMULATION AND SPECIFICATION OF THE AKF *STRIKE LIMIT ALGORITHM*

Overview

In the AKF-*SLA*, a simple population model is used to describe the stock dynamics. Furthermore, the relationship between observed stock size and true stock size is assumed to be linear. The stock dynamics model and the observation model contain a number of unknown parameters, some of which are fixed and others estimated by Bayesian methods in conjunction with the Kalman filtering estimation scheme. Each of the parameters to be estimated range over a discrete set of values giving a grid in the parameter space. A prior probability distribution is given to the parameter combinations in the grid and a Kalman filter is associated with each combination.

The probability associated with each parameter combination in the grid is updated by Bayesian methods each time a new survey estimate becomes available. The estimate of the state associated with each of the combinations is updated at the same time by the corresponding Kalman filter.

[1] Conditioning is the process of selecting the values for the parameters of the operating model so that this model is consistent with existing data for the species/stock.

Thus, for each combination in the grid, there corresponds the posterior probability of the particular combination and an estimate of the state (i.e. stock size) conditional on this combination. This amalgamation of Kalman filtering and Bayesian methodology is known as Adaptive Kalman Filtering (AKF). The overall estimate of the present state (stock size) is then obtained by summing all the stock estimates corresponding to the different parameter combinations, weighted by the respective posterior probabilities. This overall stock estimate is not used in the *SLA* described here.

Once a catch control law is specified, a strike limit can be calculated – conditional on the values of the parameters and the corresponding stock estimate. Associated with each conditional strike limit is the most recent posterior probability of the particular parameter combination. A set of strike limits (one for each parameter combination) with associated probabilities therefore results. Arranging all the strike limits in an increasing sequence, the associated probability distribution enables the construction of the cumulative probability distribution for the strike limit. Specifying a percentile of this distribution gives the eventual strike limit. The procedure on which the *SLA* is based is illustrated schematically in Fig.1.

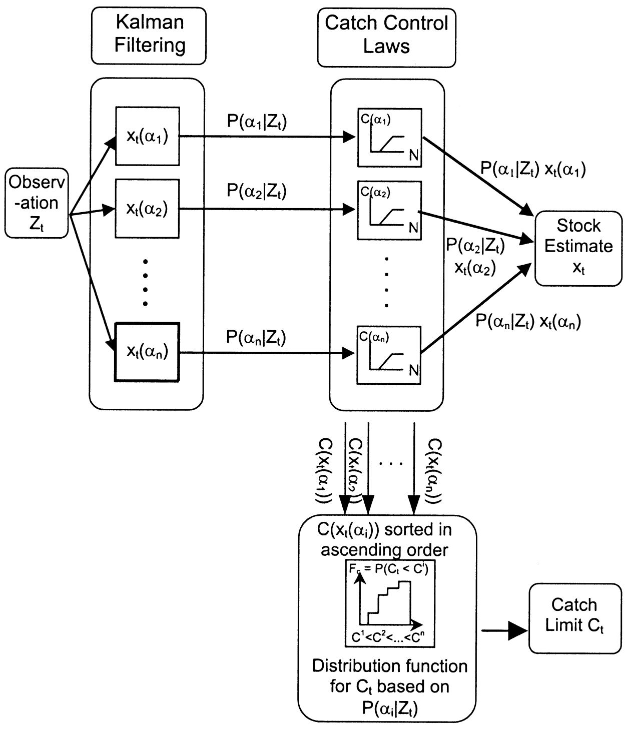

Fig. 1. An overview of the algorithm for setting the strike limit.

Some of the specifications in the AKF-*SLA* refer to the B-C-B stock, but the algorithm proposed here is nevertheless fairly general and should be applicable to a range of stocks with fairly minor modifications. The choice of the parameter grid on which the state estimate is conditional is also fairly flexible, but in this version the estimates are conditional on values of *MSYR* and *K*, the pre-exploitation population size (carrying capacity), in a two dimensional grid.

The Kalman Filter

It is assumed that the population dynamics and observations are governed by the following equations:

$$N_{t+1} = \left(S(N_t - C_t) + (1-S)\left(1 + A\left(1 - \left(\frac{N_t}{K}\right)^z\right)\right) N_t \right) e^{u_t} \quad (1)$$

$$N_t^{obs} = BN_t e^{v_t} \quad (2)$$

where N_t is the total population of animals 1 year and older (1+) in year t, C_t is the catch in year t and u_t and v_t are normal random variables with zero mean and variances Q_t and R_t, respectively. This is the well-known Pella-Tomlinson (P-T) model with parameters: annual survival rate S; pre-exploitation population size K; and the resilience parameter A, which is related to $MSYR$ by $MSYR = A\cdot(1–S)/S\cdot z/(z + 1)$. The exponent z in equation (1) determines the MSY-level ($MSYL$) according to $MSYL = (z + 1)^{1/z} K$. This model is in fact a simplification of the usual P-T models since there is no delay in the dynamics here. The possibility of biased observations is addressed by the parameter B, which is termed a bias factor. Note also that the process noise enters by simply multiplying the usual P-T function by a lognormal random variable.

The state variable is defined to be $x = ln(N)$ and the observation $y = \ln(N^{obs})$, or $y = \ln(N^{obs}) - \ln(B)$ if the possibility of bias is considered. The state and observation equations can therefore be written in the form:

$$x_{t+1} = f(x_t) + u_t \quad (3)$$

$$y_t = x_t + v_t \quad (4)$$

where

$$f(x_t) = \ln\left(S(e^{x_t} - C_t) + (1-S)\left(1 + A\left(1 - \left(\frac{e^{x_t}}{K}\right)^z\right)\right)e^{x_t}\right) \quad (5)$$

The state and observation model (3) and (4) is in a form which lends itself to state estimation via the Extended Kalman Filter (the equation describing the dynamics is non-linear and hence the EKF — in which non-linear functions are linearised – must be used). The model is of a particularly simple form; it is one-dimensional and the function relating the state and observation is simply the identity. The equations comprising the Kalman filter will therefore have a relatively simple form. The function $f(x)$ is clearly non-linear and in order to apply the Kalman filtering method a linearisation is required, i.e.

$$F_t = \frac{\partial f}{\partial x}(x) = \frac{S\cdot e^x + (1-S)\cdot e^x\left(1 + A\left(1-\left(\frac{e^x}{K}\right)^z\right) - A\cdot z\cdot\left(\frac{e^x}{K}\right)^z\right)}{S(e^x - C_t) + (1-S)\left(1 + A\left(1-\left(\frac{e^x}{K}\right)^z\right)\right)e^x} \quad (6)$$

To start a Kalman filter estimation process, an estimate of the state x at $t = 0$ needs to be specified, together with the corresponding error covariance matrix, which in this case is simply a scalar variance since the model is one-dimensional. If the population is assumed to be at carrying capacity at $t = 0$, then the initial conditions can be taken to be $x_0 = K$ and $P_0 = 0$, since there is no initial variance as K is specified in the model.

The estimate of the state at time t, using data up to t-1 is denoted by $x_{t|t-1}$ and is known as the prior estimate of x_t. The corresponding variance at time t is:

$$P_{t|t-1} = E((x_t - x_{t|t-1})^2) \quad (7)$$

When a new observation y_t becomes available, the estimate $x_{t|t-1}$ is updated according to:

$$x_{t|t} = x_{t|t-1} + K_t(y_t - x_{t|t-1}) \quad (8)$$

which is the posterior estimate of x_t i.e. the estimate of the state at time t using data up to t. Here K_t is known as the Kalman gain at time t. Note that the term in brackets on the right hand side is the difference between the actual observation and the predicted observation at time t. Thus a large difference between the actual and predicted observations will give a large modification in the state estimate and a small difference results in a correspondingly small modification. The Kalman gain is given by:

$$K_t = P_{t|t-1}(P_{t|t-1} + R_t)^{-1} \quad (9)$$

The variance $P_{t|t-1}$ is updated by:

$$P_{t|t} = (1 - K_t)P_{t|t-1} \quad (10)$$

Note that $P_{t|t}$ is the variance associated with the updated (posterior) estimate of the state at time t.

Finally, new prior estimators of the state and the error matrix at t+1, are obtained by the forward projection equations:

$$x_{t+1|t} = f(x_{t|t}) \quad (11)$$

$$P_{t+1|t} = F_t P_{t|t} F_t^T + Q_t \quad (12)$$

where F_t is given by equation (6) and the linearisation is about the point $x = x_{t|t}$. The Kalman gain at time t+1 can then be calculated and hence the posterior estimate of the state at t+1 and so on.

Equations (8)-(12) are the recursive equations for the extended discrete Kalman filter. If the model equations are linear, then the particular form of the gain K_t given by equation (9) minimises the value of $P_{t/t}$, i.e. the mean square estimation error. Note that P is simply a scalar in this model. The Kalman filter is therefore the optimal linear estimator for systems with linear observations and dynamics. For a non-linear model the question of optimality is more problematic. Note also that the effect of the observations on the updated state estimate depends on the relative values of state noise and measurement noise. The Kalman gain K, increases with the former but decreases with the latter.

Bayesian estimation of model parameters

In the state-estimation scheme described above, it is asssumed that all parameter values are known. This is seldom the case and the present model contains five unknown parameters, S, A, z, K and B, in addition to the variances Q and R. In the application described in this paper, three of those, z, S and B are fixed and the others – i.e. the resilience parameter A and the carrying capacity K – are estimated by Bayesian methods. Each of the two parameters ranges over a sequence of discrete values. This gives a 2-dimensional grid of values (A_i,K_j), $i = 1,2,\ldots,I$; $j = 1,2,\ldots,J$ with IJ different sets of parameter values. Once the parameters have been fixed, the EKF can be applied; thus to each of the IJ parameter sets there corresponds an extended Kalman filter. Whenever a new observation becomes available, the stock estimate, $x_{t|t-1}(A_i,K_j)$, is updated as described above and the posterior probability distribution, $p(A_i,K_j|Y_t)$, is updated for each of the IJ parameter sets, (A_i,K_j) by Bayesian methodology. Here, Y_t is the set of observations up to and including time t. The probability distribution is updated as follows (let α denote the vector of parameters *(A,K)*).

A prior distribution, $p(\alpha_i)$, $i = 1,2,\ldots,IJ$, for the vector α is given, and each time a new observation becomes available, a posterior distribution, $p(\alpha_i|Y_{t-1})$ is updated according to:

$$p(\alpha_i \mid Y_t) = \frac{p(Y_t \mid \alpha_i)\,p(\alpha_i)}{p(Y_t)} \tag{13}$$

where the conditional distribution $p(Y_t|\alpha_i)$, is given by the recursive formula (dropping the index on α for convenience of notation):

$$p(Y_t \mid \alpha) = \frac{1}{(2\pi)^{1/2}(P_{t|t-1}+R_t)^{1/2}} \exp\left(-\frac{(y_t - x_{t|t-1})^2}{2(P_{t|t-1}+R_t)}\right) p(Y_{t-1} \mid \alpha) \tag{14}$$

where $x_{t|t-1}$, and $P_{t|t-1}$ depend on α_i and are obtained by the Extended Kalman Filter method. Note that a 'small' prediction error $y_t - X_{t/t-1}$, gives a 'high' value of $p(Y_t|\alpha_i)$. Finally, $p(Y_t)$ is calculated by:

$$p(Y_t) = \sum_{j=1}^{IJ} p(Y_t \mid \alpha_j)\,p(\alpha_j) \tag{15}$$

This scheme for updating the state estimate and the conditional probability distribution associated with the parameter values is the Adaptive Kalman Filter (AKF). See Dereksdóttir and Magnússon (2001) for a slightly more general formulation.

Catch control law

Applying a catch control law corresponding to each of the (A_i,K_j) to $x_{t|t}(A_i,K_j)$, a total of IJ strike limits $C(x_{t|t}(A_i,K_j),A_i,K_j)$ are obtained, together with the associated posterior probability distribution $p(A_i,K_j|Y_t)$, $i = 1,2,\ldots,I$; $j = 1,2,\ldots,J$. Arranging all the IJ strike limits $C(x_{t|t}(A_i,K_j),A_i,K_j)$ in an increasing sequence, the associated probability distribution makes it possible to construct the cumulative distribution function $F(C)$ for the strike limit. Once a percentile γ of this distribution is set, a provisional strike limit is determined by solving for C_t.

$$F(C_t) = p(C < C_t) = \gamma \tag{16}$$

A one-parameter family of catch control laws is used in this version of the *SLA*. The conditional strike limit is $C = qMSY$ if the stock size N exceeds $MSYL$ and is given by the rule $C = qRY$, where RY is the replacement yield as calculated from equation (1), if $N < MSYL$. The multiplier q is a function of the conditional estimate of the stock size (i.e. conditional on K and $MSYR$) and is chosen from a family of continuous piecewise linear functions. This family is parameterised by β, the q-value at $0.5MSYL$. The multiplier q depends on N as follows

$$q = \begin{cases} 0 & N < 2000 \\ \dfrac{\beta}{(0.5MSYL - 2000)}(N - 2000) & 2000 < N < 0.5MSYL \\ \dfrac{(0.8-\beta)}{0.4MSYL}(N - 0.5MSYL) + \beta & 0.5MSYL < N < 0.9MSYL \\ \dfrac{N - 0.9MSYL}{MSYL} + 0.8 & 0.9MSYL < N < MSYL \\ 0.9 & MSYL < N \end{cases} \tag{17}$$

The strike limits as functions of stock size are shown in Fig. 2, which also shows the *RY*-curve for reference. The parameter β is a measure of the steepness of the catch control law and is one of the two parameters whereby the procedure can be tuned i.e. the level of risk chosen. The other is the percentile γ in the cumulative distribution function (cdf) for the nominal strike limit. Fig. 3 shows the cdf for four different β values and illustrates how a higher value of β gives a higher strike limit for a fixed γ-value. This difference disappears as γ approaches one. A strike limit is then set as:

$$SL_t = min(C_t, need_t) \tag{18}$$

where $need_t$ is the pre-specified level of aboriginal need in year t. All components refer to the 1+ component of the population, i.e. the total number of animals one year and older.

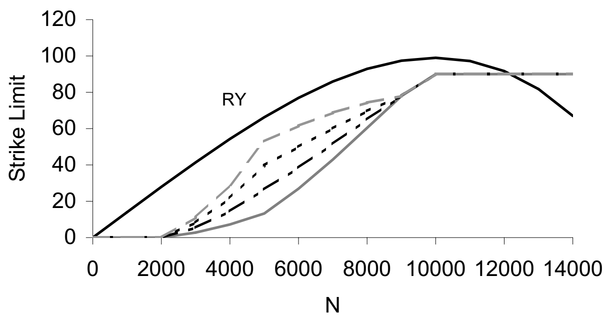

Fig. 2. Strike limits for β = 0.2, 0.4, 0.6 and 0.8 and calculated replacement yield (*RY*) as a function of stock size. *MSYL* is set to 10,000 and *MSYR* to 1%.

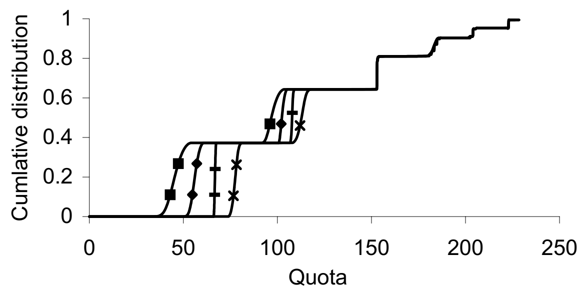

Fig. 3. The cumulative distribution function for the strike limit at the beginning of management (2003) for four values of β. Key: ■ = Beta 0.2; ◆ = Beta 0.4; — = Beta 0.6; X = Beta 0.8.

A so-called 'snap to need' feature is incorporated whereby the strike limit is increased to need if the provisional strike limit resulting from the *SLA* exceeds 95% of need, and finally, a maximum of 20% change in strike limits between years is imposed. The strike limit is set for 5-year blocks at a time.

Implementation for the B-C-B bowhead stock

A few more specifications are required. The B-C-B version described here is defined as follows. The population parameters which are kept fixed, *S* and *z*, are set to 0.99 and 2.39 respectively (this value of *z* corresponds to an *MSY*-level of 60% of *K*). The carrying capacity *K*, ranges from 10,000 to 23,000 in increments of 100 and the values of the resilience parameter *A* correspond to *MSY*-rates of 1%, 1.5%, 2%,…,3.5% and 4%. The possibility of a survey bias is not considered in this version, i.e. B is set equal to one. There are therefore 131 values of *K* and 7 values of *MSYR* in the parameter grid, giving a total of 917 filters. It is assumed that the stock is at carrying capacity in 1848 when commercial whaling began. The filters are therefore started in that year, with initial conditions $x_0 = K$ and $P_0 = 0$. The state *x* is projected forward by equation (11) and not updated until 1978 when the first survey estimate was made. On the other hand, the variance *P* is projected forward every year by equation (12). The variables *x* and *P* are updated by equations (8) and (10) respectively, whenever a new observation becomes available. There are 10 historical abundance observations between 1978 and 1993, and the *SLA* will be given abundance observations in 2002, 2004 and then every 5 or 10 years in the evaluation trials (IWC, 2003a). To each observation there is an associated estimate of the Coefficient of Variation (*CV*). The variance of the measurement noise v_t , is given by:

$$R_t = Var(v_t) = \ln(1 + CV(N_t^{obs})^2) \qquad (19)$$

In order to obtain a value for the variance of the process noise *Q*, some simulation experiments were carried out using a simplified population model to generate stock data. *Q* was chosen so as to roughly minimise prediction error – i.e. error between actual and predicted stock size − in a small subset of simulation trials. This gave $Q = 10^{-3}$. Note that a high value of *Q* gives a high Kalman gain and hence the estimates tend to follow the individual observations, which is not a desirable feature. The sensitivity to the value of *Q* was investigated and the results did not give any reason for changing the value of *Q* (Dereksdóttir and Magnússon, 2001).

Since there is no prior information on the values of the parameters *A* and *K*, the prior distribution for each parameter set (A_i, K_j), $i = 1,2,..7$; $j = 1,2,…,131$, is assumed to be discrete uniform on the specified grid. The first update of this probability distribution is made in 1978.

RESULTS

The set of *Evaluation Trials* defined by the Standing Working Group (SWG) on the Development of an Aboriginal Whaling Management Procedure (IWC, 2003a) is given in Table 1. These trials were designed to evaluate the performance of potential *SLA*s for the B-C-B stock of bowhead whales by simulating management over a period of 100 years. Factors such as *MSYR*, *MSYL*, survey interval, survey bias, survey CV and changes in need are varied and stochasticity in population dynamics introduced in order to investigate how well the *SLA*s perform under different scenarios. One hundred replicates of each trial are simulated and medians and percentiles of various performance statistics calculated from these 100 replicates. Definitions of all performance statistics are given in IWC (2003b); the key ones are the following: final depletion − relative to *K* − after 100 years of management (D1); relative recovery (D10), which is the ratio of final population size over the population size at the start of management; and need satisfaction (N9) over the first 20 years and over the entire 100-year period. A value of 1.000 means that aboriginal need was fully satisfied over the period in question. The key trials are the following:

(1) BE01, which is considered to represent the most likely picture of stock history and parameter values.
(2) BE04 where observations are negatively biased, i.e. underestimate the true abundance. Strike limits are likely to be lower than necessary and need satisfaction (N9) therefore too low. The challenge is to keep need satisfaction as high as possible.
(3) BE09, which is a low *MSYR* trial where full need satisfaction is not possible.
(4) BE12, which is an extreme trial from a conservation perspective, i.e. low *MSYR*, positively biased observations with the bias changing with time and underestimated survey *CV*. The stock is therefore likely to end up very depleted and the challenge is to prevent this as far as possible.
(5) BE16 where the increase in aboriginal need is greater than in BE12, but is a less extreme trial as regards survey quality.

Any *SLA* should in general have one or more free parameters, which determine the performance with respect to one or more criteria. These performance criteria are usually mutually incompatible and a trade-off is therefore unavoidable. For an *SLA*, the key trade-off is between risk of stock depletion and poor need satisfaction. Values in one region of parameter space make the procedure conservative in the sense that the risk of depleting the stock is small, but need satisfaction may be unnecessarily low. On the other hand, parameter values in a different region may give high need satisfaction, but at a higher risk to the stock. The parameters used for setting the level of risk are referred to as tuning parameters. The tuning of this version of the AKF-*SLA* is two-dimensional, the two parameters being β, the value of *q* at 0.5*MSYL*, see equation (17), and γ, the percentile in the cumulative distribution for conditional strike limits. The full range of values for β and γ is [0,0.8] and [0,1] respectively, with higher values of either giving higher strike limits.

Table 2 shows the results from the application of the AKF-*SLA* with what is referred to as the baseline tuning to the set of *Evaluation Trials*. The values of (β,γ) are (0.70,0.35). Results are shown for the four key performance statistics. The full set of results is given in IWC (2003b), which also gives the results for a more conservative tuning, i.e. $(\beta,\gamma) = (0.55,0.35)$. This low tuning was selected in order to demonstrate that the *SLA* can be tuned to give median relative recovery (D10) in BE12 greater than one, i.e. final stock size greater than the stock size at the beginning of management. The guidelines for the selection of the baseline tuning parameters were to maintain full median need satisfaction in trial BE01, to keep the median final depletion in BE12 above 30% of *K* and to try to maximise need satisfaction in BE04. These results will not be discussed in any detail here; suffice to say that the only ‘problem trials’ are BE04 where need satisfaction is not satisfied in spite of the stock being able to withstand such removals, and BE12 and BE13 where the recovery is not satisfactory. Furthermore, stochastic population dynamics do not appear

Table 1

The *Evaluation Trials* for the BCB stock of bowhead whales. Initial need is 67 in all cases. 10-year survey interval unless specified. Any changes in need or survey bias are linear in time (IWC, 2003a).

Trial no.	$MSYR_{1+}$	$MSYL_{1+}$	Final need	Historical survey bias	Future survey bias	Survey CV (true, estimated)		Other
BE01	2.5%	0.6	134	1	1	0.25	0.25	
BE02	2.5%	0.6	67	1	1	0.25	0.25	
BE03	2.5%	0.6	134	1	1→ 1.5 in yr 25	0.25	0.25	
BE04	2.5%	0.6	134	1	1→ 0.67 in yr 25	0.25	0.25	
BE04a	2.5%	0.6	134	1	1→ 0.67 in yr 25	0.25	0.25	5yr surveys
BE05	2.5%	0.6	134	1	1	0.25	0.25	
BE07	2.5%	0.8	134	1	1	0.25	0.25	
BE08	2.5%	0.6	134	1	1	0.25	0.25	15yr surveys
BE09	1%	0.6	134	0.67→1	1	0.25	0.25	
BE09a	1%	0.6	134	0.67→1	1	0.25	0.25	5yr surveys
BE10	4%	0.8	134	1	1	0.25	0.25	
BE10a	4%	0.8	134	1	1	0.25	0.25	5yr surveys
BE11	2.5%	0.6	134	1	1→1.5 in yr 25	0.25	0.10	
BE12	1%	0.6	134	1→1.5	1.5	0.25	0.10	
BE12a	1%	0.6	134	1→1.5	1.5	0.25	0.10	5yr surveys
BE13	1%	0.6	67	1→1.5	1.5	0.25	0.10	
BE14	2.5%	0.6	201	1	1	0.25	0.25	
BE16	1%	0.6	201	0.67→1	1	0.25	0.25	
BE20	4%	0.8	201	1	1	0.25	0.25	
BE21	U[1,4%]	0.6	134	1	1	0.25	0.25	
BE22	2.5%	0.6	134	1	1	0.25	0.25	20yr lag
BE23	2.5%	0.6	201	1	1	0.25	0.25	Strategic surveys
BE24	0.6%	0.6	134	1	1	0.25	0.25	Inertia model

Table 2

Results of the BCB *Evaluation Trials* for the AKF-*SLA* with baseline tuning. Trials labelled 'SE' have stochastic population dynamics.

	D1: Final 1+ Depletion			D10: Relative Increase			N9: Need Satisfaction (20 years)			N9: Need Satisfaction (100 years)		
	5%	Median	96%	5%	Median	96%	5%	Median	96%	5%	Median	96%
BE01	0.860	0.862	0.907	1.154	1.228	1.345	0.950	1.000	1.000	0.859	1.000	1.000
BE01-SE	0.834	0.873	0.925	1.103	1.231	1.380	0.960	1.000	1.000	0.855	1.000	1.000
BE02	0.934	0.935	0.936	1.250	1.315	1.424	1.000	1.000	1.000	1.000	1.000	1.000
BE03	0.860	0.862	0.874	1.154	1.215	1.313	0.960	1.000	1.000	0.983	1.000	1.000
BE03-SE	0.831	0.868	0.907	1.102	1.223	1.338	0.980	1.000	1.000	0.976	1.000	1.000
BE04	0.915	0.926	0.937	1.242	1.301	1.409	0.940	1.000	1.000	0.722	0.774	0.824
BE04a	0.920	0.927	0.935	1.239	1.307	1.416	0.960	1.000	1.000	0.711	0.769	0.805
BE05	0.860	0.863	0.905	1.156	1.233	1.343	0.910	1.000	1.000	0.855	0.978	1.000
BE07	0.964	0.964	0.965	1.142	1.209	1.345	0.980	1.000	1.000	0.985	1.000	1.000
BE07-SE	0.946	0.964	0.981	1.082	1.206	1.359	0.970	1.000	1.000	0.947	1.000	1.000
BE08	0.860	0.862	0.905	1.152	1.222	1.341	0.970	1.000	1.000	0.873	1.000	1.000
BE09	0.495	0.570	0.624	1.037	1.175	1.304	0.950	1.000	1.000	0.798	0.900	1.000
BE09a	0.514	0.568	0.625	1.055	1.182	1.277	0.960	1.000	1.000	0.814	0.890	1.000
BE09-SE	0.499	0.578	0.688	0.987	1.152	1.336	0.950	1.000	1.000	0.795	0.901	1.000
BE10	0.976	0.985	0.988	0.997	1.010	1.028	0.970	1.000	1.000	0.817	0.930	1.000
BE10a	0.976	0.985	0.988	0.997	1.012	1.028	0.990	1.000	1.000	0.826	0.930	1.000
BE11	0.861	0.863	0.884	1.154	1.218	1.317	0.930	1.000	1.000	0.958	0.993	1.000
BE12	0.241	0.317	0.393	0.746	0.948	1.135	0.870	1.000	1.000	0.738	0.821	0.961
BE12a	0.273	0.321	0.388	0.798	0.969	1.108	0.860	1.000	1.000	0.738	0.821	0.907
BE12-SE	0.224	0.326	0.428	0.710	0.991	1.217	0.920	1.000	1.000	0.734	0.824	0.969
BE13	0.361	0.425	0.510	1.176	1.286	1.375	0.950	1.000	1.000	0.980	1.000	1.000
BE14	0.798	0.838	0.905	1.082	1.192	1.344	0.900	1.000	1.000	0.699	0.896	0.968
BE16	0.389	0.538	0.622	0.789	1.107	1.307	0.890	0.980	1.000	0.630	0.757	0.958
BE16-SE	0.377	0.543	0.645	0.784	1.069	1.324	0.880	0.980	1.000	0.632	0.758	0.952
BE20	0.967	0.982	0.988	0.993	1.009	1.028	0.910	1.000	1.000	0.657	0.854	0.969
BE21	0.720	0.938	0.974	1.025	1.225	1.554	0.960	1.000	1.000	0.847	1.000	1.000
BE22	0.842	0.894	0.942	1.316	1.396	1.481	1.000	1.000	1.000	1.000	1.000	1.000
BE22-SE	0.805	0.884	0.999	1.188	1.366	1.563	1.000	1.000	1.000	0.966	1.000	1.000
BE23	0.780	0.833	0.899	1.072	1.166	1.329	0.920	0.980	1.000	0.705	0.915	0.976
BE24	0.549	1.256	1.382	1.342	2.622	2.977	0.950	1.000	1.000	0.815	0.969	1.000

to have much effect; if anything, the results are slightly better. The survey interval is not crucial either in this particular set of trials, although the result is marginally better in BE12 with a shorter interval.

In order to explore how performance depends on the tuning parameters, 20 (β,γ) combinations in the square [0.5,0.8]$\times$[0.2,0.6] were selected. Note however, that these are not the most extreme tunings possible. Fig. 4 shows plots of the BE01-N9, BE04-N9 and BE12-D10 statistics for the full 100-year horizon. This figure demonstrates the trade-off between the recovery statistic BE12-D10 and the need satisfaction statistics, median of BE04-N9 (Fig. 4a) and 5th percentile of BE01-N9 (Fig. 4b), as well as the relationship between the two need satisfaction statistics (Fig. 4c). The 5th percentile of N9 in BE01 is selected rather than the median, because the latter hits the upper bound of 1.00 for a large fraction of the (β,γ) combinations. It is clear that some (β,γ) combinations are dominated, i.e. it is possible to find other combinations where performance is superior in both need satisfaction and recovery. An indication of the location of the optimality frontier, i.e. where improvement in one statistic is not possible without a deterioration in the other, can be seen in Fig. 4a. This frontier illustrates well the trade-off between recovery as measured by BE12-D10 and need satisfaction as measured by BE04-N9. However, this frontier is far from being optimal for BE12-D10 and BE01-N9(5%) as is clear from Fig. 4b. The selection of values for tuning parameters, even if a level of a recovery or depletion is fixed, is thus very difficult. For example, suppose that the condition BE12-D10 = 0.95 is imposed. Then the (β,γ) combination (0.7,0.3) is clearly superior to (0.5,0.6) in terms of BE04-N9, but the reverse is the case for BE01-N9(5%).

The SWG also specified a large number of other trials, named *Robustness Trials*. These trials include such factors as time-varying biological parameters, episodic events, where two events occur in years 1-50, in which 20% of the animals die, etc. The full specifications of the *Robustness Trials* and the results for the AKF-*SLA* are given in IWC (2003b) and will not be discussed here, with the exception of a trial in which the population crashes to 2,000 whales in the first year, but everything else is as in trial BE16. This scenario may not be very plausible, but serves a useful purpose in assessing how quickly the *SLA* reacts to drastic changes. The resulting values for the D1 statistic (final depletion) with 10-year survey intervals is 0.000; 0.000; 0.044 for 5th percentile, median and 95th percentile. It is therefore quite clear that the *SLA* does not react fast enough to prevent extinction in the majority of the replicates. One reason is the 20% maximum change between years, which is imposed on the strike limit. Decreasing the survey interval to five years gives D1 statistics 0.060; 0.119; 0.161 which is a considerable improvement. This can be improved even further by removing the 20% bound on changes in strike limit. The higher survey frequency is therefore crucial in preventing extinction in such a scenario.

Fulfilling aboriginal need is a high priority goal. This is the motivation for the 'snap-to-need' feature where the strike limit is increased to need if the provisional strike limit resulting from the *SLA* is close to need, i.e. greater than a certain percentage thereof. Need increases with time in most trials. 'Snapping-to-need' has therefore less of an impact on the stock size early in the management period, since the resulting increase in strike limits in terms of numbers of whales is lower. It may therefore be argued that the percentage could safely be set lower initially, thus increasing need satisfaction in the first years of management. In order to investigate the effect of a variable 'snap-to-need' percentage, a few trials were run in which, for the first 20 years, the strike limit was set to need if the *SLA* strike limit exceeded 80% of need and, thereafter, set to need if the *SLA* strike limit exceeded 95% of need. The results are given in Table 3. Need satisfaction in the first 20 years is almost fully satisfied without any significant deterioration in depletion statistics for trials such as BE09, BE12 and BE16. The improvement in need satisfaction in years 1-20 is further illustrated in Table 4 which gives the number of replicates where need is not fully satisfied in this period. A lower initial 'snap-to-need' percentage is thus a definite improvement.

DISCUSSION

The *SLA* described in this paper rests on a sound theoretical basis, i.e. the well-established Kalman filtering techniques and Bayesian methods. The state estimation part of the *SLA* employs the so-called Extended Kalman Filter – which is applicable to non-linear systems – to estimate the stock size conditional on (*MSYR*,*K*) values. Bayesian methodologies are used for calculating strike limits. The only part of the

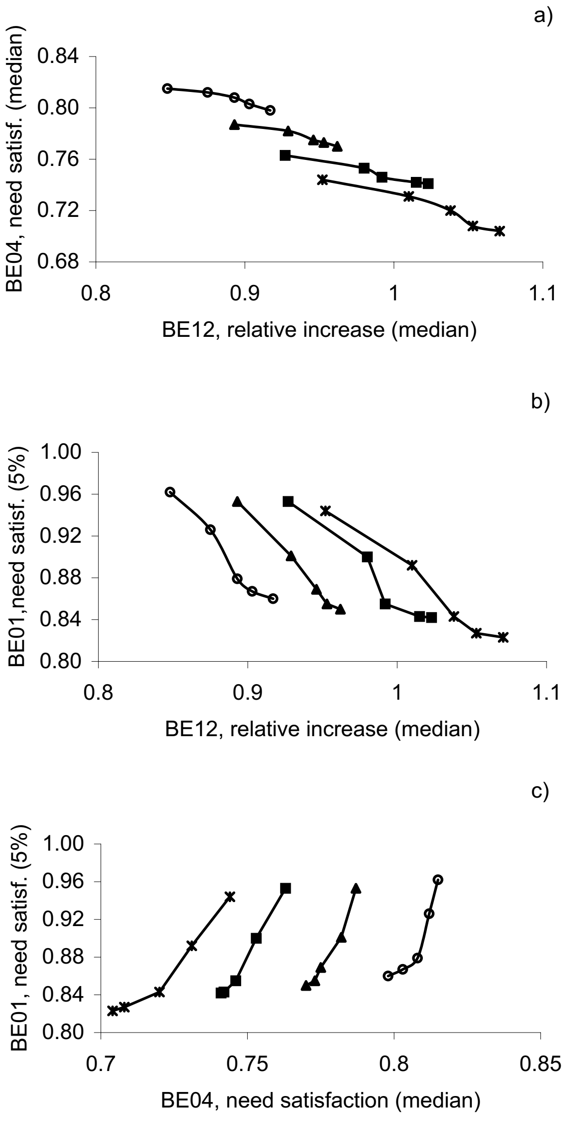

Fig. 4. Median need satisfaction in BE04, relative increases in BE12 and need satisfaction in BE01 for three values of β and γ from 0.2 to 0.6 in increments of 0.1. In (a) and (b) γ decreases from left to right for each β value and in (c) γ increases from left to right for each β. Key: ✱ = Beta 0.5; ■ = Beta 0.6; ▲ = Beta 0.7; ● = Beta 0.8.

Table 3

The results for baseline tuning of the AKF-*SLA* when snap-to-need is lowered to 80% the first 20 years and kept at 95% thereafter.

	D1: Final 1+ Depletion			D10: Relative Increase			N9: Need Satisfaction (20 years)			N9: Need Satisfaction (100 years)		
	5%	Median	96%	5%	Median	96%	5%	Median	96%	5%	Median	96%
BE01	0.860	0.862	0.907	1.154	1.228	1.344	1.000	1.000	1.000	0.859	1.000	1.000
BE04	0.915	0.926	0.937	1.242	1.301	1.409	1.000	1.000	1.000	0.724	0.778	0.828
BE05	0.860	0.863	0.905	1.156	1.233	1.341	1.000	1.000	1.000	0.862	0.982	1.000
BE09	0.495	0.570	0.622	1.037	1.164	1.300	1.000	1.000	1.000	0.799	0.911	1.000
BE10	0.976	0.985	0.988	0.997	1.010	1.028	1.000	1.000	1.000	0.818	0.931	1.000
BE12	0.241	0.315	0.390	0.746	0.943	1.135	0.946	1.000	1.000	0.738	0.822	0.964
BE12a	0.268	0.318	0.383	0.798	0.961	1.108	0.941	1.000	1.000	0.745	0.821	0.907
BE14	0.798	0.837	0.903	1.082	1.189	1.324	0.942	1.000	1.000	0.700	0.908	0.968
BE16	0.389	0.532	0.617	0.789	1.097	1.285	0.939	1.000	1.000	0.647	0.764	0.958
BE20	0.967	0.981	0.988	0.993	1.009	1.028	0.949	1.000	1.000	0.671	0.854	0.969
BE21	0.720	0.938	0.974	1.025	1.225	1.554	1.000	1.000	1.000	0.847	1.000	1.000
BE23	0.788	0.837	0.900	1.072	1.176	1.328	1.000	1.000	1.000	0.709	0.906	0.977

Table 4

Number of replicates where need is not fulfilled in the first 20 years. Results for the baseline case are shown in the column on the left and results for the baseline case with snap-to-need lowered to 80% the first 20 years is shown in the column on the right.

	Snap-to-need fixed at 95%	Snap-to-need lowered to 80% the first 20 years
BE01	19	0
BE04	27	0
BE05	29	2
BE09	19	0
BE10	8	0
BE12	42	7
BE12a	47	7
BE14	46	8
BE16	52	10
BE20	36	5
BE21	10	0
BE23	52	1

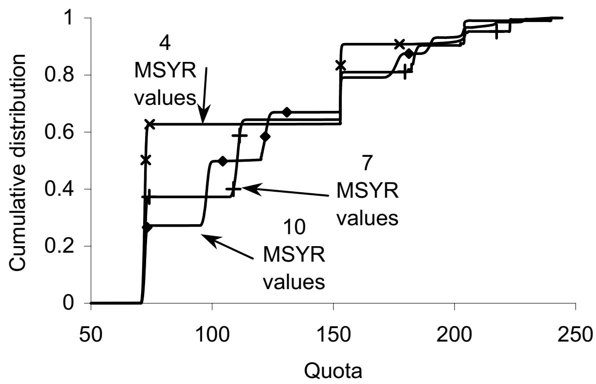

Fig. 5. The cumulative distribution function for the strike limit at the beginning of management (2003) for three sets of (*MSYR*, *K*) filters.

SLA which could be regarded as *ad hoc* is the set of catch control laws, but these laws are nevertheless of a general form used by the IWC in the past.

The *SLA* is tuned by two independent parameters giving substantial tuning flexibility and a fairly wide range of results. There is however some scope for improvement.

It is clear from Fig. 3 that the cumulative distribution function for the strike limit is close to a step function with the steps corresponding to the different *MSYR* values used in the filters. This is reflected in plots of catch trajectories, where distinct bands corresponding to the *MSYR* values are apparent in a few trials (IWC, 2003b). It is therefore desirable to reduce the height of the steps and make the strike limit change more smoothly with the tuning parameter γ. This can easily be achieved by increasing the number of *MSYR* filters. Fig. 5 shows the cumulative distribution function at the beginning of management (2003) for 4, 7 and 10 *MSYR* filters. The height of the steps is lowest for 10 *MSYR* filters as expected. Ideally, the grid for *MSYR* values should be as fine as possible, but for computational reasons 'only' 7 filters are used in this version of the AKF-*SLA*. Another way to smooth the cumulative distribution function is to fit polynomials or spline functions through the step function, but this has not been explored here.

The possibility of using a set of bias filters, i.e. to condition on parameter values in a 3-dimensional grid was not explored here. This is investigated in Dereksdóttir and Magnússon (2001), where it is shown that the AKF is very successful in identifying a negative survey bias. Performance in such a trial, BE04, is thus greatly improved, i.e. need satisfaction increased substantially. On the other hand, performance deteriorated somewhat in the conservation trials (BE09 and BE12) since the AKF was less successful in identifying the positive bias. There are nevertheless good reasons for keeping the option of bias filters open, when applying the AKF-*SLA* to other stocks.

The overall stock estimate is not used in this version of the AKF-*SLA*. It may be worthwhile to explore the possibility of using this estimate or the corresponding time-trends therein to change the values of the tuning parameters during the simulation period. Similarly the overall stock estimate may be used to invoke a protection level. There is no explicit protection level in the present version, only an implicit one coming from the catch control law.

Some further investigations of the properties of this *SLA* are desirable. For example, it is of interest to see how the weighted average of *MSYR* develops as more and more observations become available and to compare this estimate to the true *MSYR* value. It is also of interest to see how well the overall stock estimate − obtained as the weighted average of the individual estimates (see Fig. 1) − tracks the true stock trajectory. These aspects were investigated to some extent in Dereksdóttir and Magnússon (2001) where it was shown that this estimate tracked the true stock fairly well and that the estimation part of the *SLA* was indeed learning as more and more observation became available.

REFERENCES

Brown, R.G. and Hwang, P.Y.C. 1997. *Introduction to Random Signals and Applied Kalman Filtering*. 3rd Edn. Wiley, New York. 484pp.

Dereksdóttir, E.H. and Magnússon, K.G. 2001. Investigations of an Aboriginal Whaling Management Procedure using Adaptive Kalman Filtering (AKF). *J. Cetacean Res. Manage.* 3(2):109-16.

Donovan, G.P. (ed.). 1982. *Reports of the International Whaling Commission (Special Issue 4). Aboriginal/Subsistence Whaling (with special reference to the Alaska and Greenland fisheries)*. International Whaling Commission, Cambridge, UK. 86pp.

Gudmundsson, G. 1994. Time series analysis of catch-at-age observations. *Appl. Stat.* 43(1):117-26.

Hammond, P.S. and Donovan, G. (eds.). In press. *Special Issue on the Revised Management Procedure*. International Whaling Commission, Cambridge, UK.

International Whaling Commission. 1999. Chairman's Report of the Fiftieth Annual Meeting. *Ann. Rep. Int. Whaling Comm.* 1998:3-61.

International Whaling Commission. 2003a. Report of the Fourth Workshop on the Development of an Aboriginal Subsistence Whaling Management Procedure (AWMP). *J. Cetacean Res. Manage. (Suppl.)* 5:489-499.

International Whaling Commission. 2003b. Report of the Scientific Committee. Annex E. Report of the Standing Working Group on the Development of an Aboriginal Subsistence Whaling Management Procedure (AWMP). *J. Cetacean Res. Manage. (Suppl.)* 5:154-225.

Kalman, R.E. 1960. A new approach to linear filtering and prediction problems. *J. Basic Eng.* 82:34-45.

Kimura, D.K., Balsiger, J.W. and Ito, J.H. 1996. Kalman filtering the delay-difference equation: practical approaches and simulations. *Fish. Bull.* 94:678-91.

Pella, J.J. 1993. Utility of structural time series models and the Kalman filter for predicting consequences of fishery actions. pp. 571-94. *In:* G. Kruse, D.M. Eggers, R.J. Marasco, C. Pautzke and T.J. Quinn II (eds.) *Proceedings of the International Symposium on Management Strategies for Exploited Fish Populations*. Alaska Sea Grant College Program Report 93-02, University of Alaska, Fairbanks. 836pp.

Reed, W.J. and Simons, C.L. 1996. Analysing catch-effort data by means of the Kalman filter. *Can. J. Fish. Aquat. Sci.* 53:2,157-66.

Stoker, S.W. and Krupnik, I.I. 1993. Subsistence whaling. pp. 579-629. *In:* J.J. Burns, J.J. Montague and C.J. Cowles (eds.) *Special Publication*. No. 2. *The Bowhead Whale*. 1st. Edn. Society for Marine Mammalogy, Lawrence, Kansas.

Date received: June 2002
Date accepted: October 2002.