W1: Modelling and Monitoring Environmental Outcomes in Adaptive Management

at iEMSS summit, Burlington, Vermont, July 2006.

John Norton[a] and Ken Reckhow[b]

[a] *iCAM and Dept. of Mathematics, Australian National University, Canberra*

[b]*Nicholas School of the Environment and Earth Sciences, Duke University, Durham, North Carolina*

Abstract
The main aim of this paper is to stimulate questions about the future of adaptive management (AM) of natural resources, and more specifically about what approaches may be feasible which have not yet been explored well. The method will be to compare the histories, strengths and limitations of AM, control engineering and Bayesian analysis, which have superficial similarities, significant differences and perhaps lessons for each other. Some difficulties encountered by control engineering and complex environmental models are pointed out.

Keywords: *adaptive management, control engineering, Bayesian analysis, feedback, modelling, monitoring*

1. INTRODUCTION: ADAPTIVE MANAGEMENT AND FEEDBACK CONTROL

Adaptive management (AM) [Holling, 1978] is an approach to environmental management which explicitly confronts future uncertainty due to prediction model deficiencies, variability in modelled inputs, unpredictable and/or unmodelled disturbances, unexpected and unmodelled outcomes, and changes in objectives and priorities among interested parties. The main means suggested for tackling these difficulties are:

(1) designing management as a continuing trial-and-error learning procedure, in which some variation in the state of the system is accepted as valuable because it yields more information about the system's behaviour;

(2) comparison, through selected indicators, of the results of alternative management policies, rather than attempted optimisation of some cost function;

(3) inclusion of resilience in the face of disturbance as an objective; and

(4) emphasis on the importance of monitoring and remedial mechanisms.

These features have some apparently quite close parallels in control engineering, which is also concerned with trying to regulate or cause desired changes in systems which may be complex and not very well modelled, and subject to unmeasured, unpredictable disturbances. The ostensible similarities are worth closer examination to see how far they really extend and hence what, if anything, the evolution of control engineering may suggest about the future of natural resource management. The similarities and essential differences between the two fields may also put the difficulties faced by environmental management in a new light. To simplify the discussion, we shall be thinking mainly about water-resource management, although ecological sustainability is the application considered in much of the AM literature.

Figure 1 shows the essential components of a feedback control system. Some analogies with AM can be drawn. "Plant" corresponds to the environmental system to be managed, "reference input(s)" to desired environmental outcome(s), "sensors" to measurement (or at least assessment) of environmental indicators, "control synthesis" to the whole process of reaching a management decision from desired outcomes and monitored behaviour, and "actuators" to the translation of management decisions into action on the ground. The plant is typically modelled imprecisely and selectively, the reference inputs tell only part of the story about what is desired, the sensors are limited in what they can measure and they introduce error (systematic and random), the control synthesis can be tackled in many different ways, and the actuators impose hard constraints on the location, nature, extent and speed of action. Alongside these rough correspondences are others, no less significant but not in the picture, similarities between the motivations, problems, tactics and constraints. The following sections discuss them, first in general terms then with reference to perceived successes and failures in the past 25 years of AM and 75 years of feedback control. Remaining unresolved difficulties are stressed and some questions are posed about implications for AM.

2. SHARED AND DISTINCT FEATURES OF THE MANAGEMENT AND CONTROL PROBLEMS

It would be naïve to expect engineering methodology to provide solutions to complex environmental management problems, for many reasons, most obviously the often large uncertainties, the prominence of the social-institutional dimension in the latter, and the sometimes ill defined scope of the problem. That said, the contexts of environmental management and engineering design have similarities strong enough to warrant comparison. Engineering design is sometimes perceived as dealing with accurately known components and systems and strictly formulated and relatively simple performance criteria. This is usually far from the truth. Much of engineering design is about reducing sensitivity of performance to variations in components, systems and their working environments. Performance criteria are a mixture of often conflicting factors such as cheapness, reliability and safety, technical performance, ease of maintenance, and subtler aspects such as attractiveness to the client, familiarity, novelty and distinctiveness. An example is an aircraft jet engine: a large, complex, distributed, time-varying, dynamical system which must perform reliably and economically over a wide range of air pressure and temperature, Mach number and thrust, and must not enter conditions such as surge and rotating stall.

In control engineering and the earlier use of feedback in communications subsystems, the original and still by far the most important motivation for the application of feedback [Maciejowski, 1989; Bennett, 1996] is to reduce the sensitivity of the overall relation between desired and actual output(s) to variations in the plant and to disturbances. The sensitivity-reduction property of feedback is readily seen by reference to Figure 2. The plant P is controlled by action generated by a controller C operating on the difference E between the reference input (desired output) R and the result FY of feeding back the measured output Y (ignoring measurement error for now) along the feedback path F. P, F, and C can be thought of as operators (gains if they are linear, transfer functions if they have linear dynamics). The overall behaviour is given by

$$Y = PCE = PC(R - FY) \text{ so } Y = \frac{PC}{1+PCF}R$$

and a plant variation δP causes an output change

$$\delta Y = \left(\frac{P+\delta P}{1+(P+\delta P)CF} - \frac{P}{1+PCF} \right) CR$$

so the sensitivity, defined as the ratio of proportional changes in output and plant, is

$$\frac{\delta Y / Y}{\delta P / P} = \frac{1}{1+(P+\delta P)CF}$$

Clearly this is small if $|1+(P+\delta P)CF| >> 1$, *i.e.*, if the perturbed "loop gain" $|(P+\delta P)CF|$ is much greater than 1. Without feedback the sensitivity would be 1. (Although this sketch of an analysis applies to linear, single-input, single-output systems, analogous properties apply to linear, multivariable and some well behaved non-linear systems.)

In addition, the output component due to an additive disturbance input D at the plant output is $Y_D = \frac{1}{1+PCF}D$, reduced by feedback under exactly the same conditions as sensitivity to plant changes. This relation could alternatively be expressed in terms of "stiffness" $\frac{D}{Y_D} = 1 + PCF$: the higher the loop gain, the larger the disturbance needed to cause any given output change.

These properties of feedback systems (subject to limitations discussed later) link straightforwardly with item (4) of the list of AM features, as measurement and feedback form a "monitoring and remedial mechanism". Attenuation of the effects of disturbances on the output is not quite the same as asking for resilience (item (3) in AM). Resilience is the ability not to be pushed into the domain of attraction of a different and worse equilibrium, and is thus a matter of controlling the state and possibly inputs to stay within some region, rather than reducing excursions of the output. There is a large body of control theory concerned with maintaining stability of non-linear systems when perturbed by external disturbances or internal variations.

The connection of the classical control configuration of Figure 1 with item (1) of AM is more complicated. The development of all but the simplest control systems is largely by trial and error, i.e. learning by doing, aided by models, as recommended for AM. Whether the control scheme itself operates by trial and error is less clear. In feedback control, the corrective action applied to the plant is error-actuated, but a scheme such as that in Figure 1 with *fixed F* and *C* would not be considered adaptive. Adaptive control is discussed later.

The "monitoring and remedial mechanism" provided by measurement and feedback has undergone extension and generalisation in control engineering, for reasons with some relevance to AM. The first part of the feedback process, measurement, introduces error ("observation noise"). Attenuation of this error should appear among the objectives of the control/management system, as above. Measurement also has to contend with incomplete accessibility of the variables which one would like to monitor. The standard solution in control engineering is to design an observer or state estimator which reconstructs estimates of the inaccessible variables of interest. The state observer/estimator employs a model of how the state variables influence the accessible outputs, directly or through other state variables [Maybeck, 1979, 1982; O'Reilly, 1983]. The observer may be designed to have "optimal" probabilistic properties such as unbiasedness and statistical efficiency or orthogonality (uncorrelatedness) between estimates and errors, or simply to attenuate measurement errors and to converge at an acceptable rate to the correct values in the absence of errors. Some free parameters in the observer are tuned to get a satisfactory compromise between speed of convergence and sensitivity to error. The observer is itself a predictor-corrector, using the discrepancy between observed and predicted output values to modify the estimates of the inaccessible variables. In other words, the observer also incorporates feedback. ***Is there scope in AM to make monitoring more systematic by thinking of it as state estimation?***

An indication of how far the idea of an observer might be useful in AM can be gained by seeing what an observer depends on. It requires identification of a set of state variables, which must meet two conditions. The state variables must be self-contained, in the sense that their values at one instant and the values of all the forcing inputs affecting them from then onwards determine (ideally) their later values. They must also be observable in the formal sense that the measured variables, accompanied by a model relating the state and measured variables, are enough to determine fully the initial, and thence any later, state. The model relating forcing, state and observed variables need not be accurate, as model-induced error will be reduced to some degree by the corrector part of the predictor-corrector, but of course must not be very misleading. A bigger problem is that the unmodelled error in the measurements must not be so structured as to be mistaken for systematic behaviour of the measured variables. ***Questions for AM are whether there is any possibility of model-based state estimation, and whether systematic measurement errors can be separated from the underlying "signal".***

Item (2) of AM advocates that the outcomes of alternative management policies should be compared rather than attempting optimisation. Control design is mostly a matter of finding a solution which meets a number of requirements expressible as inequalities, e.g. on speed of response to a step change in desired output, or on the margins by which stability is maintained. Any remaining design freedom is used to optimise a performance or robustness measure, but optimality is not the first concern. Over a period of about 15 years from about 1960, control design was largely pitched as variational optimisation, finding the control policy yielding the best state trajectory or minimising a terminal cost (such as the time to reach a desired condition or the final error between the output values and their desired values). This optimal control approach, ignoring uncertainty, works well with good models and in predictable environments, as in spacecraft control, but has largely been abandoned elsewhere, for two reasons. First, optimal designs are often found not to be robust to modelling error. Second, the designer has to specify weights in the cost function to be optimised, e.g. weighting control effort against state error. Finding suitable weights is thus a trial-and-error process of examining performance against weights, so the apparent simplicity of a once-and-for-all, hands-off optimisation is illusory. ***This piece of history can be taken to support the view that AM should not be approached simply as an optimisation problem.*** However, there may be scope for borrowing *robust* optimisation-based methodology such as model predictive control, discussed later.

3. ADAPTIVITY

3.1 Limitations of feedback and motivation for adaptivity

There are severe restrictions on the use of feedback to reduce sensitivity to plant variation and disturbances. First and best known, high loop gain may (depending on the dynamics of the system) be limited by instability. Roughly speaking, instability arises when the gain is high enough and the delay right for a disturbance to be reinforced as it passes round the loop, through the plant, sensor, feedback, error formation and controller. A common erroneous intuition is that negative feedback (as in Figure 1) implies stability. This fails to consider the combined effects of gain, delay in passing round the loop and sign reversal in forming the error. The combined effect may be that an oscillation that has gone round the loop returns as a larger, in-phase, reinforcing version of itself. If the plant, measurement and actuator characteristics are known fairly well, the controller can be designed to avoid this condition and ensure stability. However, a trade-off usually has to be made between the margin by which instability is avoided and the degree of attenuation of disturbances and measurement error. Typically the need to avoid instability limits the range of speeds (or equivalently frequencies) over which good attenuation is achievable. Similar comments apply to multi-input, multi-output systems and to some types of non-linear systems.

Thus in design of feedback control systems, consideration of the time scale of the desired response and of disturbances and errors is critical, to recognise the restrictions imposed by stability and because speed of response (as well as reduction of slowly varying error, i.e. regulation of the steady state) is part of the design requirement.

For AM with monitoring and remedial mechanisms, i.e. feedback, the same risk of instability applies. In particular, delay in the system itself or in measuring indicators and revising management action may have ill effects, ranging from overshoots or undershoots (due to reduced stability margin) to runaway (instability). Action that would be beneficial if timely can easily be damaging if delayed. Moreover, awareness that short- and long-term water-resource management concerns may differ has begun to gain exposure [Fath and Beck, 2005]. One can conclude that consideration of time scale of response should also be prominent in AM of environmental systems. ***Is it?***

An additional problem for AM is that uncertainty or changes in the system and/or its inputs may be enough to make a controller unexpectedly unstable, even if the model says it should not be. If we design the feedback conservatively, to keep the system stable over the whole range of expected behaviour, the performance will usually be much poorer than if our model and assumed inputs were exact. One approach to dealing with poorly predictable behaviour is to try to make the controller (or AM scheme) adapt to whatever behaviour of the plant (or environmental system) eventuates. The name of adaptive management supposes some such ability. The history of adaptive control will now be reviewed very briefly to see if it has any lessons for AM.

3.2 Adaptive control and its failings

Definitions of "adaptive" vary in control engineering but agree in referring to the controller having its parameter values changed as conditions are found to change. The mere presence of feedback, correcting control action on the basis of what the output is seen to do, does not imply adaptivity. If "learning by doing" in AM encompasses changing management *policy* as well as changing management *action* in response to the results of monitoring, it is adaptive control.

The history of adaptive control in engineering is quite long [Åström and Wittenmark, 1995; Åström, 1996; Bennett, 1996, Bushnell, 1996] and salutary. Successive heavy bursts of research have resulted mainly in realisation that practical adaptive control is liable to a number of severe problems. Model-reference adaptive control, the favoured approach in the 1950s and 1960s, mainly in aerospace applications, tried to make the response of the controlled system follow a specified ideal, model response, in the face of variation in forcing and plant response (e.g. aerodynamics varying with speed and altitude). The need to alter the controller parameters imposed a risk that stability would be lost. Lyapunov theory [Åström and Wittenmark, 1995] provides a way to analyse stability of an unforced, non-linear system and see what will ensure it. However, as the stability characteristics of a non-linear system generally depend on the forcing as well as the parameters, the danger remains.

Dual control [Feldbaum, 1961] aims to use the control signal to carry out two tasks: keeping the output well behaved while probing the system to improve knowledge of its dynamics. The idea of deliberate perturbation experiments, i.e. probing, is also an intrinsic part of AM. It clearly requires some compromise to avoid too much disturbance of the system. In the event, analytical design of dual control to achieve the best compromise still proves too complicated for all but the simplest examples.

An alternative is to identify a parametric model of the plant, updated regularly, and employ it in synthesising a control law which will thus follow any changes in the plant behaviour. This is called self-tuning control [Åström and Wittenmark, 1973; Wellstead and Zarrop, 1991]. With a suitably simple model structure, parameter-estimation technique and control objective, such a scheme can be quite economical, especially if the option is taken of identifying and updating the controller parameters directly. Furthermore some stability analysis is possible. However, it was discovered that the system tends to become too confident in the model parameter values and fails to react to sudden changes. If the system is kept alert to change and none occurs, there is a danger that the system gradually becomes prone to excessive reaction, bursting at long intervals into violent oscillation then recovering [Anderson, 1985; Åström and Wittenmark, 1995]. Work on cures is still current [Apley, 2004].

A conclusion from this history is that allowing the parameters of a controller to vary as conditions change is dangerous, for two reasons. The parameter adjustment is based on assumptions about the nature of plant changes or disturbances, which if infringed may lead to collapse of the system. Secondly, analysis of a closed-loop system with performance-linked changes in model or controller parameters is as yet only feasible for simple, usually over-idealised systems, so unexpected and unintended behaviour is possible and performance guarantees are not available. These facts have motivated strong emphasis over the last 25 years on robust control [Green and Limebeer, 1995; Zhou and Doyle, 1998; Ackermann, 2002]. Robust control aims to get a performance guarantee valid while the plant and disturbances remain in specified ranges. A heuristic yet widely successful approach to robust control, Model Predictive Control, is outlined in the next section.

Updating of the model and the action taken is fundamental in both AM and feedback control. Bayesian analysis offers a broad yet prescriptive framework for updating, so its possible relevance to AM will also be discussed.

4. PROBLEMS IN ADAPTIVE MANAGEMENT AND SOME TOOLS FROM OTHER FIELDS

4.1 A short list of problems in adaptive management

A concise list of problems encountered in trying to apply AM in ecological management is given by Johnson [1999]: "...difficulties in developing acceptable predictive models, conflicts regarding ecological values and management goals, inadequate attention to nonscientific information, and an unwillingness by agencies to implement long-term policies seen as too risky or costly". These have counterparts in control engineering, but with differences which may provoke thought. They are discussed in turn below.

4.2 "Difficulties in developing acceptable predictive models"

Difficulties in developing acceptable predictive models are common also in the process industries, because of process complexity and non-linearity, variations due to changes in ambient conditions and feedstock and ageing effects, and the distributed nature of much process plant. Two lines have been taken to mitigate this lack of good models. The first is to rely on local control, using simple controllers to regulate individual variables on the assumption that good local control everywhere implies good control overall, so long as local objectives are well chosen. In other words, the overall problem is split into a hierarchy of levels, with objectives for local, short-term control set at a higher level and varied on a longer time scale, with higher levels still of longer-term control for production scheduling, economic optimisation and maintenance.

The second tactic to deal with lack of a good predictive model is to avoid relying on accurate long-term prediction, but to exploit the ability of even a poor model to give fairly good guidance to short-term results of control actions and some indication of long-term results. This is the basis of Model Predictive Control (MPC) [Garcia et al., 1989; Morari and Zafiriou, 1989; Soeterboek, 1992; Bemporad and Morari, 1999; Kouvaritakis and Cannon, 2001; Maciejowski, 2002].

4.3 Robustness to poor prediction via Model Predictive Control

MPC evolved in the petrochemical industry to handle problems of controlling poorly modelled, complex processes subject to stringent operating constraints. Its basis is constrained optimisation, model-based, over a fairly long future period, with the resulting control policy applied for a short while but then reoptimised to take account of the measured response of the plant to this initial action. One advantage of such a scheme is that even if the predictive model is not very good, tentative optimisation taking the relatively long-term future into account makes the initial part of the control actions more circumspect than would optimisation over a short term. Another advantage is that it is an "open-loop feedback" scheme, with straightforward optimisation treating the plant as not under feedback control, but the resulting control periodically corrected by the fed back information in the measurements. Jointly these advantages confer robustness, i.e. only gradual degradation of performance by deviations of actual behaviour from predictions. ***Are any of the features of MPC likely to be applicable and effective in AM?***

MPC is not a panacea, though. An excessively misleading prediction model will cause breakdown, as will too-large unforeseeable disturbances or cumulative modelling error. Some increase in robustness may be attainable if (a big "if") the uncertainties can be characterised in advance, e.g. by the probabilistic properties of abrupt disturbances or the correlation structure of more coherent disturbances, or by bounds on modelling error. Recently Carlson and Doyle [2000] have argued that designs highly optimized and robust in one range of assumed circumstances can be fragile and fail catastrophically when the design assumptions are infringed; robustness implies fragility. This worry has a counterpart in the recognition [Gunderson and Holling, 2002] that an ecological system may go abruptly from variations within an acceptable range of conditions to much worse conditions, i.e. from around a "good" equilibrium to near a bad one. AM reacts by emphasising high resilience as an aim, where resilience is the size of disturbance which can be withstood without leaving the domain of attraction of the "good" equilibrium point. Most robust control engineering (but not MPC) reacts by adopting minimax design, minimising deterioration (maximising performance) in the worst case over a prescribed range of plant uncertainty and, in some cases, disturbances. That approach is open to the objection that average performance is likely to be poor if an unlikely worst case is catered for. Rather surprisingly, little academic robust control research has tried to overcome the objection by finding ways to gauge the situation well enough to alter the control objective according to how far the plant is from disaster. ***Would minimax policy objectives make sense in AM?***

A related question arising both control engineering and AM is how to deal with asymmetrical performance criteria. For instance, in many (perhaps most) cases a deviation from the desired output value in one direction is more serious than an equal deviation in the other. Similarly the significance of errors in a state estimate may be asymmetrical. For analytical convenience symmetrical weighting is often assumed (*e.g.* minimising mean-square or mean absolute error). However, the increasing trend to numerical solution of control problems confers flexibility in how errors are weighted.

4.4 Adaptive Management and Bayesian Analysis

Effective updating on the basis of new information is central to adaptive management. Bayesian analysis should be particularly useful for adaptive management because of the natural way in which knowledge can be updated as new information becomes available, via Bayes' Theorem. Bayes' Theorem can be written as:

$$\pi(\theta|y) = \frac{\pi(\theta)\, f(y|\theta)}{\int_\theta \pi(\theta)\, f(y|\theta)\, d\theta}$$

where $\pi(\theta|y)$ is the probability (density) of the value θ of the parameter vector after observing the new data, y, (i.e. the posterior probability of θ), $\pi(\theta)$ is the probability of θ before observing y (the prior probability of θ), and $f(y|\theta)$ is the likelihood function which incorporates statistical relationships as well as the mechanistic or process relationships among the predictor and response variables. As θ is integrated out of the denominator of the expression for $\pi(\theta|y)$, this simple, logical expression stipulates that, when combining information, the resultant (or posterior) probability is proportional to the product of the probability reflecting a priori knowledge (the prior probability) and the probability representing newly acquired knowledge (the sample information, or likelihood), since the denominator does not depend on θ.

Information synthesis is usually the motivation for employing Bayesian analysis; thus Bayesian analysis serves as an excellent approach for the analytics of adaptive management. The conventional application of a Bayesian approach emphasizes the combination of prior information and a single set of data (post-implementation monitoring data). However, it is shown in Bayesian statistics texts that sequential updating, using the posterior from the previous step as prior for each successive step, is equivalent to updating using all of the data together; thus sequential updating provides a means to investigate possible temporal patterns in the data, which can be attractive for adaptive management.

Sequential Bayesian updating also provides an appealing interpretation of the standard predictor-corrector state estimator, the Kalman filter [Maybeck, 1979, 1982]. Using θ to denote state for the moment, we wish to update the conditional mean (unbiased, minimum-mean-square-error) estimate $\hat{\theta}_{k-1|k-1}$ of system state at time k-1, based on information Y_{k-1} up to and including the observations making up vector $\mathbf{y}_{k-1}$, on receiving $\mathbf{y}_k$. This yields $\hat{\theta}_{k|k}$ based on Y_k. The updating consists of prediction (using $\hat{\theta}_{k-1|k-1}$ in the state equation, setting any unknown forcing inputs to their mean values), to give the conditional mean $\hat{\theta}_{k|k-1}$ of the probability density function (pdf) $p(\theta_k \mid Y_{k-1})$ and the corresponding covariance, followed by correction according to $\mathbf{y}_k$, using Bayes' rule for the probability density functions:

$$p(\theta_k \mid Y_k) \equiv p(\theta_k \mid \mathbf{y}_k, Y_{k-1}) = \frac{p(\mathbf{y}_k \mid \theta_k) p(\theta_k \mid Y_{k-1})}{p(\mathbf{y}_k \mid Y_{k-1})}$$

Here $p(\mathbf{y}_k \mid \theta_k)$ follows directly from the pdf of the observation error. As $p(\mathbf{y}_k \mid Y_{k-1})$ is not a function of θ_k, it serves only to scale $p(\theta_k \mid Y_k)$, so the conditional mean $\hat{\theta}_{k|k}$ and associated covariance can be computed without evaluating it. If all the pdfs are assumed to be Gaussian, the correction (observation-update) equations of the Kalman filter are readily derived from this relation. [The Gaussian assumption was not needed in Kalman's original derivation, which made the state-estimation error orthogonal to the estimate]. Analytical updating from $p(\theta_{k-1} \mid Y_{k-1})$ via $p(\theta_k \mid Y_{k-1})$ to $p(\theta_k \mid Y_k)$ is feasible only in a few special cases like this, but is increasingly possible for other pdfs by Monte Carlo updating of a large set of state samples from $p(\theta_{k-1} \mid Y_{k-1})$, time-updating each through the state equation then weighting it with the observation likelihood $p(\mathbf{y}_k \mid \theta_k)$.

As an example of sequential Bayesian updating for adaptive management in a US water quality standards compliance program (the USEPA TMDL Program), a series of computer programs were developed to automate the process of updating water quality concentration estimation from model predictions and subsequent monitoring data. These programs use Bayesian analysis results for (log) normal random variables, and the conjugate family of prior distributions. The process has three steps. First, a number of procedures were developed for converting model forecasts of water quality concentrations to a prior distribution of the underlying concentration distribution parameters. Second, a program was developed to produce the posterior distribution of the underlying concentration distribution parameters and the posterior predictive distribution of future observations, based on the pre-implementation model forecast (the "prior") and the first year of post-implementation monitoring data (the "sample"). Third, the "posterior" distribution of the underlying concentration distribution parameters is then converted to a prior distribution of the same parameters for the next time period, and the process repeats when new data are available.

To demonstrate this process, a Bayesian SPARROW [McMahon et al. 2003, Qian et al. 2005] model-predicted 1992 nitrogen concentration distribution for the Neuse River Estuary (North Carolina) was used to develop a prior distribution of the mean and variance of log nitrogen concentrations; the sequentially updated posterior predictive distributions for each subsequent year are presented in Figures 3 and 4. The same process was repeated for the chlorophyll a concentration distribution in the Neuse River Estuary (Figure 5). The prior distribution for chlorophyll a was developed using an empirical model [NeuBERN; Borsuk et al., 2003, 2004a,b] and the results from the SPARROW model. Although the prior distribution based on NeuBERN over-estimated the chlorophyll a concentration, the sequentially updated posterior predictive distributions (based on post-implementation monitoring data) quickly converged to a distribution similar to the observed chlorophyll a concentration data (Figure 5).

A model is a summary of our understanding about the system under study. In adaptive management, a model should be able to be updated as we accumulate more knowledge about the system. This knowledge accumulation, when reflected in the model, may be represented in terms of refined model parameter estimates, additional modules, or a new model all together. ***Are there effective ways to select the appropriate adjustments? What factors limit the suitability of Bayesian Analysis for this purpose?***

Several paradigms are available for refining parameter estimates as new information is obtained. Bayesian estimation is an appealing one, as outlined above. A second is the idea that parameter estimators are predictor-correctors, adjusting the old estimate by an amount proportional to the error it produces in predicting the observed variables (much as in state estimation). A third idea, which unifies many parameter-estimation algorithms, is that the parameter estimates should minimise the sum, over all measurements of the system output, of some non-decreasing function of the prediction errors yielded by the model [Ljung, 1978, 1999]. It seems likely that AM could make more use of these ideas, which are applicable to a wide range of model structures. Statistical techniques for choosing how many parameters to employ are also well developed, as mentioned above. However, guidelines for adding, removing or modifying sections of models of interest for AM are lacking. Perhaps what is needed is a merger of sensitivity analysis and model development. The big question is how general it could be made. ***Are the requirements, constraints and technical features of models for AM so case-dependent that no such synthesis is feasible?***

A closely related question is whether AM can offer a systematic way to distinguish modelling error due to non-ideal parameter values or deliberate reduction of model complexity (especially through aggregation, e.g. in second-order, linear, constant-coefficient rainfall-runoff models) and more serious error due to omission of an important cause-effect link (i.e. failing to include all essential state variables), or inclusion of a mistaken one. The risk of such a problem increases as the system modelled becomes more broadly defined and harder to monitor, being higher, for instance, for an ecological than for a hydrological model.

4.5 "Conflicts regarding ecological values and management goals"

For conflicts regarding ecological values and management goals, the second item cited by Johnson [1999] as an impediment to AM, the resemblance of AM to control engineering is smaller. The control-design problem is handed over to the engineer, who can make unopposed judgements on priorities (subject to usually strict limitations on cost, time and performance). Consultation is in two phases, among the clients then among the engineers, with tendering and perhaps a feasibility study as the matching-up procedure. ***Is such a separation of interest groups desirable or feasible in AM?*** The folk wisdom is that extensive consultation is crucial, but the contrary view that it encourages disagreement is beginning to be heard. Changes in objectives are also often quoted as a problem in AM; is it realistic to ask all parties to agree once and for all on the objectives for one phase of a management exercise, in the knowledge that they can be revised at a specified later stage? That would be "open-loop feedback" control of objectives.

4.6 "Inadequate attention to nonscientific information"

Inadequate attention to nonscientific information [Johnson, 1999], which may be paraphrased as ignoring informal (and perhaps mistaken) collateral knowledge, is taken care of in engineering by relying on adequate experience among the analysts and designers, and by prototyping to discover what could not be foreseen and what was wrong. AM's "learning by doing" [Walters and Holling, 1990] sounds just like prototyping. However, Bayesian Analysis provides a means for incorporation of nonscientific information into modelling and for prediction to be based on all relevant information. Use of collateral knowledge in Bayesian Analysis becomes possible through expression of this knowledge in probabilistic terms, as a "prior probability." In a general sense, the prior probability reflects that which is known before "learning by doing" begins; that is, it may reflect prior expert judgment, prior data, or some combination.

One particularly flexible Bayesian modelling approach is the Bayesian Belief Network (BBN) or Bayes Net. BBNs refer to a network of nodes and arcs characterizing marginal and conditional probabilities among the variables of interest. These probabilities are typically derived from expert judgment and data. Used in AM, the knowledge obtained through learning by doing is propagated throughout the BBN using a set of rules [Pearl, 1988]. The result is a BBN characterizing posterior probabilities, reflecting the new knowledge. ***Is there scope for incorporating BBNs in AM, keeping them updated as conditions evolve rather than confining their use to initial, once-and-for-all modelling?***

4.7 "Unwillingness by agencies to implement long-term policies"

The problem, also quoted by Johnson [1999], of an unwillingness by agencies to implement long-term policies seen as too risky or costly, points to an essential difference in time scale between AM and most, but not quite all, engineering applications of control. It is obvious that the performance of a management or control system cannot be properly assessed over less time than it takes the system to undergo a representative sample of plant variations, disturbances and, where the ambient or plant conditions are changing slowly and systematically, drift. For a natural resource system, this will usually be from a few years to a few decades. There are "risky and costly" engineering systems the performance of which must be judged on this sort of time scale, e.g. electric power generation and transmission systems, but even for them the control performance can normally be evaluated much more rapidly. There are established means of reducing exposure to long-term risk, by hedging (keeping open a "fall-back" position and not adopting one policy or design too widely until it is proven), phasing of projects with decisions to proceed or stop taken in stages, and competition among suppliers who will make their own judgements of risk and cost. ***How far can these means be applied in AM?***

A different interpretation of "hedging" (here interpreted to mean "risk avoidance") is the "margin of safety (or MOS)" required in the USEPA program (designed to achieve compliance with water quality standards). The MOS hedges the required pollutant reduction in the direction of additional water quality improvement in order to avoid the risk of further pollution. The magnitude of this hedge is supposed to be a function of the uncertainty in the forecast pollutant load reduction. Thus, the greater the forecast uncertainty, the greater is the chance that the MOS will result in overprotection and wasted resources. AM could change this prospect by allowing smaller hedges, with the learning process requiring more stringent pollutant controls when found necessary.

5. OPEN CHALLENGES FOR ADAPTIVE MANAGEMENT

5.1 Characterisation of uncertainty

Bayesian Analysis is naturally suited for adaptive management, utilizing the language of probability for updating knowledge with new information. Indeed, under the label of "adaptive implementation" (learning while doing), AM has recently been recommended by the US National Research Council [NRC, 2001] for use in the assessment of solutions for water quality standard violations (the USEPA TMDL program) in recognition of the substantial forecasting errors associated with water quality models.

One requirement of Bayesian Analysis is that uncertainties be fully quantified (and expressed probabilistically). While this is feasible for relatively simple empirical models, it remains a challenge for detailed process simulation models. For example, limited data often result in identification problems in water quality models. As a consequence, detailed water quality process models are typically parameterized using trial and error judgment, rather than an optimization technique. However, trial and error judgment generally ignores the often-significant correlations between parameters, and thus fails to provide a basis for complete error propagation.

A more general condition for large simulation models, called "equifinality" by Beven [2001], is that, for a given model specification, many different parameter values yield essentially equally good fits to the data. While this condition is a result of fitting a model with too many degrees of freedom to limited data/knowledge, Bayesian analysis in such cases does provide the opportunity for characterizing a multivariate posterior distribution reflecting near-redundancy among the parameters through large variances and covariances. However, for large process models, the error in the model specification is often difficult to isolate from parameter error due to deficiencies in the data, so a complete characterization of prior model forecast uncertainty remains a technical challenge.

5.2 Matching the model to system characteristics

Walters [1997] reviews obstacles to the effective implementation of AM. Among them he criticises reductionist modelling of the detailed processes contributing to overall behaviour (of ecological systems, in his illustrations) when an empirical model based on experiments looking at the overall input-output relations would be enough. A similar factor applies in modelling for control design. The complexity of a control system tends to be comparable with that of the model on which its design was based, so there is a high premium on keeping the model simple so as to generate a simple, fully comprehensible control system. On the other hand, it may be necessary to model the internal relations to permit controllability and observability analysis and to ensure that the design does not produce excessive excursions in internal variables.

A further reason for not relying exclusively on input-output models is that a simple non-linearity within a system can give rise to complex input-output behaviour, hard to describe let alone understand. Deterministic chaos is a well known example, but there are many others. In considering models which display some of the internal workings, the state-space paradigm which for 45 years has been so successful in control engineering does not seem to have wide currency in models for AM. "State" and "state variables" are frequently mentioned, but with little sign that they are rigorously interpreted as a minimal collection of variables with the properties sketched in the paragraph on state observers above. There are some good excuses for not employing state-variable models in environmental applications, not least the fact that a distributed system has an infinite number of them, but the process of deciding what variables are necessary to describe a situation in proper (for the application) spatial and temporal detail is more systematic and informative if regarded as constructing a set of state variables.

Choice of state variables is inseparable from deciding on scales and sampling intervals in space and time, and hence from how much lumping is warranted and what the effective bandwidths of the inputs and outputs are. Although AM stresses the need for planned experiments, the need to consider bandwidth (or power spectra or correlation functions) and aliasing [Gabel and Roberts, 1987] does not seem to have a high profile, in contrast to control engineering where it is recognised as crucial. Perhaps the reason is that the benefits of understanding "signal" characteristics and dangers of misinterpreting behaviour through aliasing in periodically sampled data pale beside the likely problems of unmodelled phenomena, poor historical records, dubious prediction in scenarios not representative of the past and guesstimated "data" supplied by experts. Even if this is so, it makes no sense to incur yet more model deficiencies by inattention to data properties and spatio-temporal scales of dynamics.

5.3 Bottom-up and top-down modelling

Walters also criticises modelling-for-management projects which get sidetracked into ever more detailed and comprehensive modelling, justified by the spurious argument that more modelling effort and detail ensures better model-based prediction. Resource and time constraints in engineering almost always force modelling to be sharply focused on the design problem at hand. More detailed modelling is undertaken only to resolve problems arising in prototype trials. ***Are there fundamental reasons why the same should not apply in AM?*** Perhaps one such reason is the wish to avoid risking serious consequences of trial actions based on initial models.

Walters takes the view that only empirical experience, not physical principles or modelling, can tell how far averaging and selecting of contributory processes is permissible. Sensitivity assessment [Saltelli et al., 2000] of a tentative model can also play a part by showing what the dominant behaviour depends critically on and what can be omitted as having only minor effects. Model reduction [Antoulas, 2005] and model-structure testing are long-standing research topics in control engineering [Veres, 1991], but can often be avoided by insisting on an initially simple model, only added to when it is essential.

6. CONCLUSIONS

As control engineering confronts some of the same problems, and employs some of the same strategies, as AM, its history raises quite a few questions for future development of AM, as summarised in bold italics above. One of the clearest lessons of control engineering is that it is feedback, not adaptivity, which is crucial in obtaining robust performance in the face of poorly predicted or unpredictable disturbances to and variations in the system to be controlled. Feedback schemes include those, like Model Predictive Control, where planned control actions over a period are modified as model deficiencies and/or unpredicted system changes become apparent.

Bayesian Analysis provides a rigorous and logical learning model for adaptive management under uncertainty. If relevant knowledge and data are expressed in probabilistic terms, then Bayes Theorem can be used to combine new information resulting from learning by doing. For complex systems, Bayesian Belief Networks can be used to model the probabilistic relationships among variables. Newly acquired data, from monitoring or experimentation based on learning by doing, can then be used to revise the BBN probabilities; this provides the analytic support for adaptive management.

A possible objection to the whole idea of comparing other fields with AM of natural resources is that the economic, social, institutional and perceptual aspects of NRM problems are both dominant and not amenable to analysis. A possible response is that problems in other fields (engineering being one, economic management another, medicine a third) also have economic, social, institutional and perceptual aspects, and that problems in those fields have been mitigated, if not always solved, by aiming to make the problem-solving procedures more systematic, consistent and open to scrutiny. A start has been made on tackling the social aspects of criteria-setting for water-resource management problems by trying to classify stakeholders [Fath and Beck, 2005], so as to understand and presumably eventually prioritise differing demands resulting from various groups' differing views of the world. ***Can AM aspire to further progress in those directions?***

REFERENCES

Ackermann, J., Robust Control: The Parameter Space Approach, 2nd ed., *Communications and Control Engineering* Series, Springer-Verlag, Berlin, 2002.

Anderson, B.D.O., Adaptive systems, lack of persistency of excitation and bursting phenomena, *Automatica, 21*, 247-258, 1985.

Antoulas, A.C., Approximation of Large-Scale Dynamical Systems, SIAM, Philadelphia, PA, 2005.

Apley, D.W., A cautious minimum-variance controller with ARIMA disturbances, *IIE Trans*. *36*, 417-432, 2004.

Åström, K.J., Adaptive control around 1960, *IEEE Control Systems Magazine, 16*(3), 44-49, 1996.

Åström, K.J., and B. Wittenmark, On self-tuning regulators, *Automatica, 9*, 185-199, 1973.

Åström, K.J., and B. Wittenmark, Adaptive Control, 2nd ed., Prentice Hall, Englewood Cliffs, NJ, 1995.

Bemporad, A. and M. Morari, Robust model predictive control: a survey, Robustness in Identification and Control, A. Garulli, A. Tesi, and A. Vicino, (eds.), *LNCIS 245*, Springer-Verlag, London, 207-226, 1999.

Bennett, S. A brief history of automatic control, *IEEE Control Systems Magazine, 16*(3), 17-25, 1996.

Beven, K., How far can we go in distributed hydrological modelling?, *Hydrology and Earth System Sciences, 5*(1), 1-12, 2001.

Borsuk, M.E., C.A. Stow, and K.H. Reckhow, An integrated approach to TMDL development for the Neuse River Estuary using a Bayesian probability network model (Neu-BERN), *Journal of Water Resources Planning and Management, 129,* 271-282, 2003.

Borsuk, M.E., C.A. Stow, and K.H. Reckhow, A Bayesian network of eutrophication models for synthesis, prediction, and uncertainty analysis, *Ecological Modelling, 173,* 219-239, 2004a.

Borsuk, M.E., C.A. Stow, and K.H. Reckhow, The confounding effect of flow on estuarine response to nitrogen loading, *J. Environ. Eng., 130,* 605-614, 2004b.

Bushnell, L.G., (ed.), Special Issue of *IEEE Control Systems Magazine* on the History of Control, *16*(3), June 1996.

Carlson, J.M., and J. Doyle, Highly optimized tolerance: robustness and design in complex systems, *Phys. Rev. Lett. 84*(11), 2529-2532, 2000.

Dorf, R.C., and R.H. Bishop, Modern Control Systems, 10th ed., Prentice Hall, Englewood Cliffs, NJ, 2005.

Fath, B.D., and M.B. Beck, Elucidating public perceptions of environmental behavior: a case study of Lake Lanier, *Environmental Modelling & Software 20*(5), 485-498, 2005.

Feldbaum, A.A., Dual control theory, Parts I and II, *Automation and Remote Control, 21*(9), 874-880, and *21*(11), 1033-1039, 1961.

Gabel, R.A., and R.A. Roberts, Signals and Linear Systems, 3rd ed., Wiley, New York, 1987.

Garcia, C., D.M. Prett, and M. Morari, Model Predictive Control-theory and practice-a survey, *Automatica, 25,* 335-348, 1989.

Green, M., and D.J.N. Limebeer, Linear Robust Control, Prentice Hall, Englewood Cliffs, NJ, 1995.

Gunderson, L.H., and C.S. Holling, Panarchy: Understanding Transformations in Human and Natural Systems, Island Press, Washington, DC, 2002.

Holling, C.S., (ed.), Adaptive Environmental Assessment and Management, John Wiley & Sons, New York, 1978.

Johnson, B.L., Introduction to the special feature: adaptive management - scientifically sound, socially challenged?, *Conservation Ecology* [online], *3*(1), 10, URL: http://www.consecol.org/vol3/iss1/art10/, 1999.

Kouvaritakis, B., and M. Cannon, (eds.), Nonlinear Predictive Control: Theory and Practice, IEEE Control Engineering Book Series, *61*, 2001.

Ljung, L., Convergence analysis of parametric identification methods, *IEEE Trans. on Autom. Control*, AC-23, 770-783, 1978.

Ljung, L., System Identification - Theory For the User, 2nd ed., Prentice Hall, Upper Saddle River, NJ, 1999.

Maciejowski, J.M., Predictive Control with Constraints, Prentice Hall, Harlow, England, 2002.

Maciejowski, J.M., Multivariable Feedback Design, Addison-Wesley, New York, 1989.

Maybeck, P.S., Stochastic Models, Estimation and Control, vol. 1, Academic Press, New York, 1979.

Maybeck, P.S., Stochastic Models, Estimation and Control, vol. 2, Academic Press, New York, 1982.

McMahon, G., R.B. Alexander, and S. Qian, Support of Total Maximum Daily Load Programs Using Spatially Referenced Regression Models, *Journal of Water Resources Planning and Management, 129,* 315-329, 2003.

Morari, M., and E. Zafiriou, Robust Process Control, Prentice-Hall Inc., Englewood Cliffs, NJ, 1989.

National Research Council (NRC), Assessing the TMDL Approach to Water Quality Management, National Academy Press, Washington, DC, 2001.

O'Reilly, J., *Observers for Linear Systems*, Academic Press, London, New York, 1983.

Pearl, J., Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference, Cambridge University Press, Cambridge, UK, 1988.

Qian, S.S., K.H. Reckhow, J. Zhai, and G. McMahon, Nonlinear regression modeling of nutrient loads in streams: A Bayesian approach. *Water Resources Research*, July 2005.

Saltelli, A., K. Chan, and E.M. Scott, (eds.), Sensitivity Analysis, Wiley, Chichester, UK, 2000.

Soeterboek, R., Predictive Control: A Unified Approach, Prentice Hall International (UK), Hemel Hempstead, 1992.

Veres, S.M., Structure Selection of Stochastic Dynamic Systems, Gordon and Breach, New York, 1991.

Walters, C., Challenges in adaptive management of riparian and coastal ecosystems, *Conservation Ecology* [online], *1*(2), *1*, URL: http://www.consecol.org/vol1/iss2/art1/, 1997.

Walters, C.J., and C. S. Holling, Large-scale management experiments and learning by doing, *Ecology 71*, 2060-2068, 1990.

Wellstead, P.E., and M. B. Zarrop, Self-Tuning Systems: Control and Signal Processing, Wiley, Chichester, England, 1991.

Zhou, K., and J. C. Doyle, Essentials of Robust Control, Prentice Hall, Englewood Cliffs, NJ, 1998.

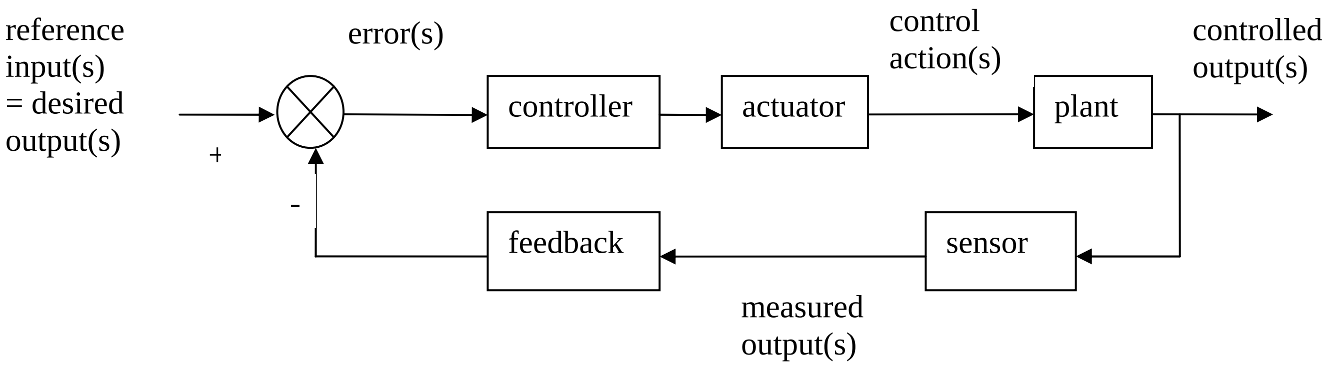

Figure 1. Feedback control system (single-variable or multivariable).

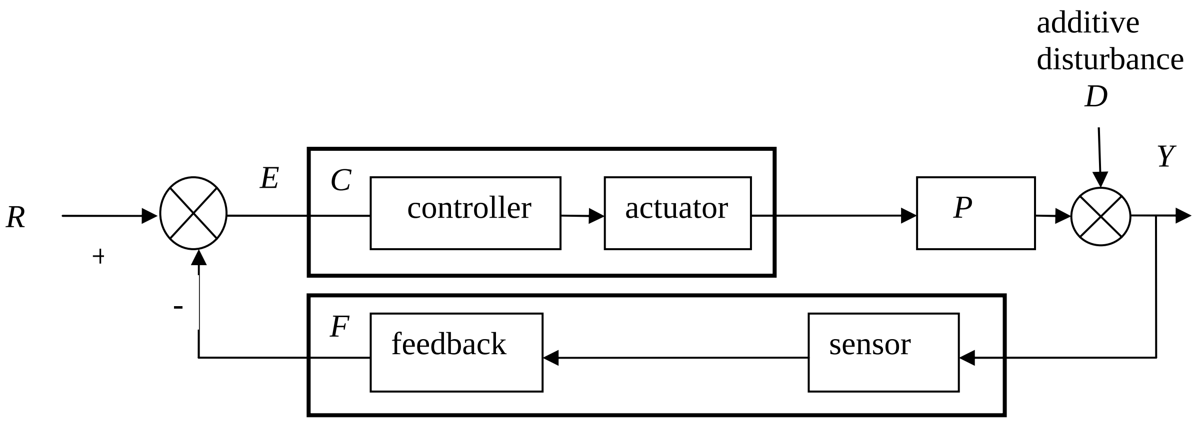

Figure 2. Feedback control system: F = sensor and feedback combined, C = controller and actuator combined.

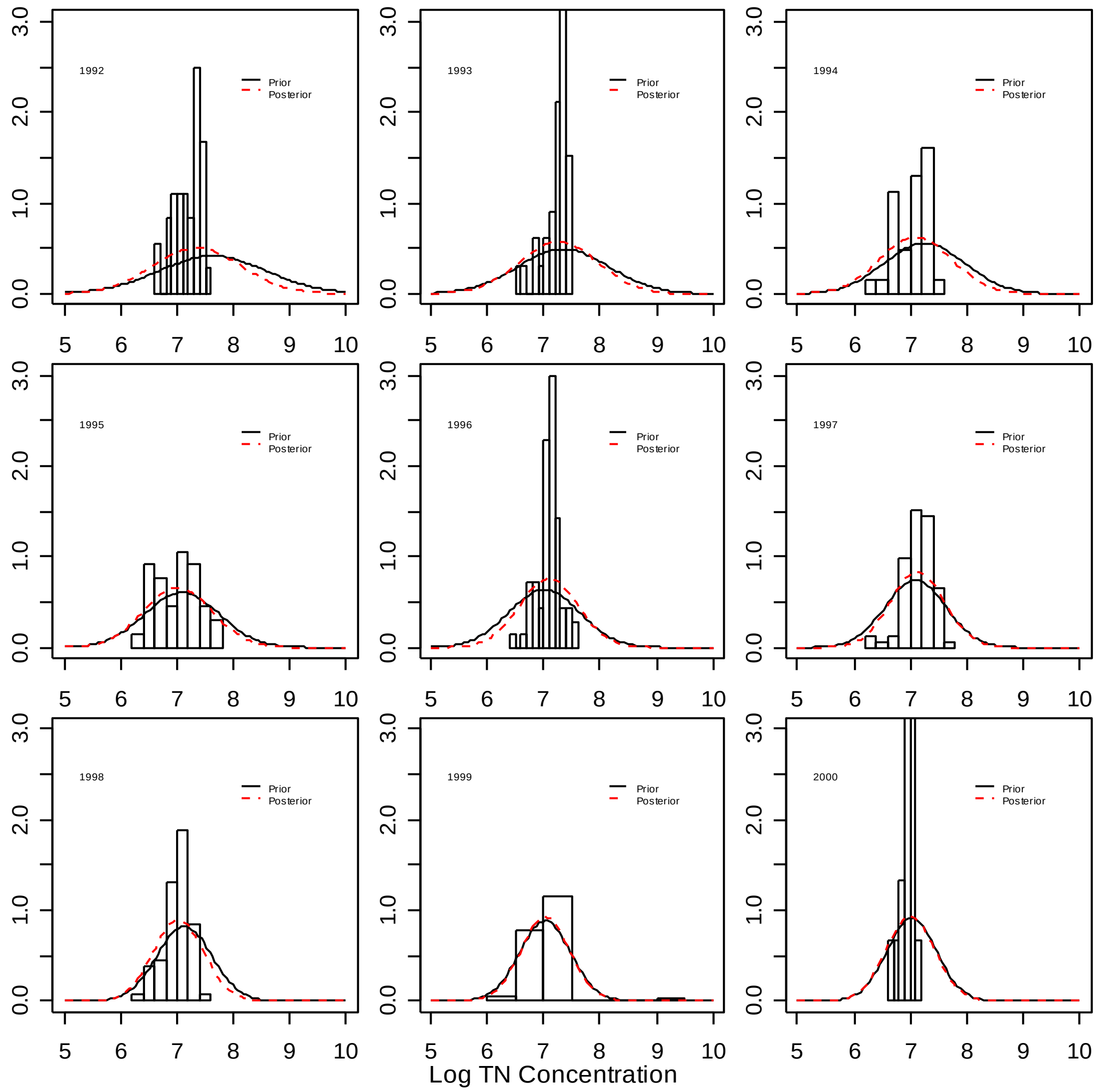

Figure 3. Sequentially updated predictive distribution of nitrogen concentration in the Neuse River Estuary. The Solid black lines are the prior distribution used for each year, the red dashed lines are the resulting posterior predictive distributions for the same year, and the data are shown in histograms.

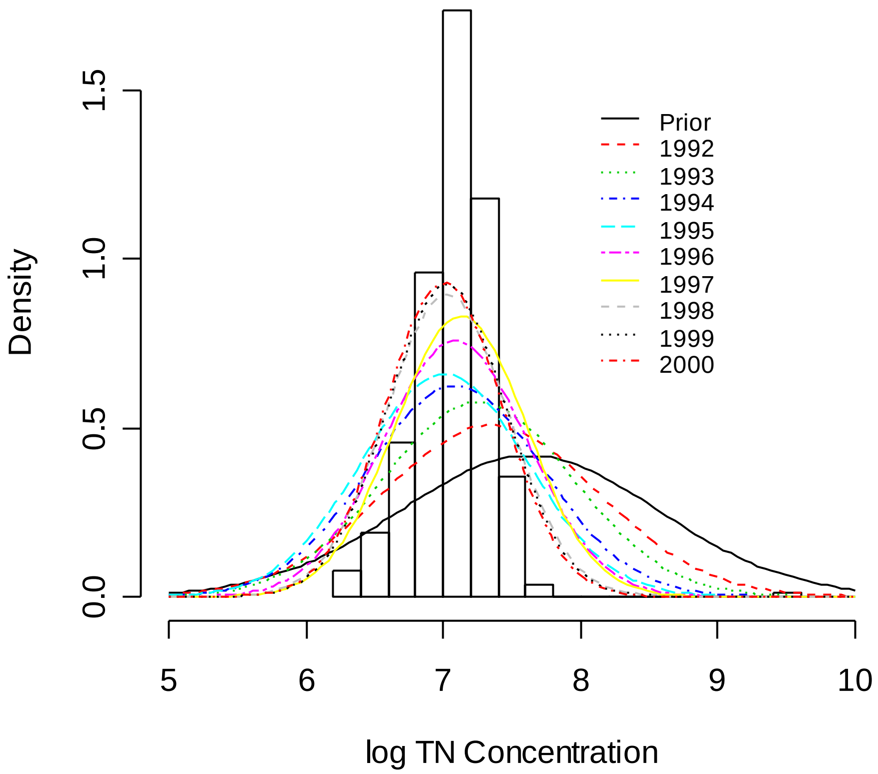

Figure 4. Sequentially updated posterior predictive distribution of log nitrogen concentrations in the Neuse River Estuary. The histogram shows the combined nitrogen monitoring data collected from 1992 to 2000.

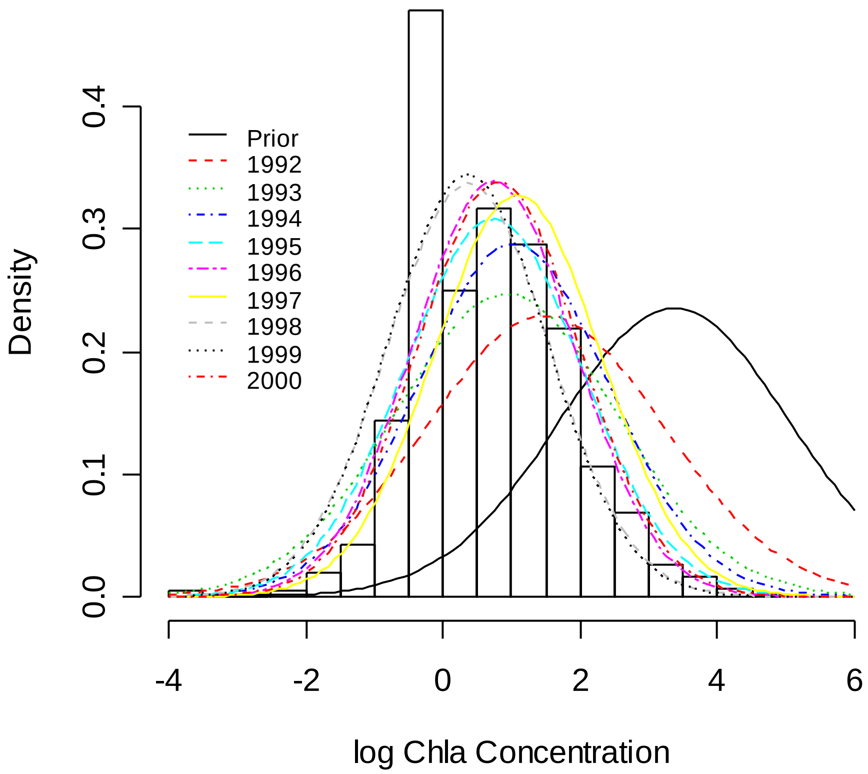

Figure 5. Sequentially updated posterior predictive distribution of log chlorophyll a concentrations in the Neuse River Estuary. The histogram show the combined chlorophyll a monitoring data collected from 1992 to 2000.