KALMAN FILTERING IN WATER QUALITY MODELING: THEORY VS. PRACTICE

B. P. Schrader
New England Power Service Company
Westborough, Massachusetts

S. F. Moore
Resource Management Associates
Lafayette, California

ABSTRACT

Kalman filtering is a statistical technique for computing minimum-uncertainty estimates of the states in linear uncertain dynamic systems. Filtering theory is qualitatively described. Potential applications to water quality modeling are discussed for state and parameter estimation, and model identification. Literature is reviewed on filtering applications to modeling receiving water quality. Several characteristics of environmental modeling problems are identified which may limit the filter's applicability. A case study is presented of optimal filtering applied to hydrothermal model development for a coastal power plant discharge. A simple model structure is proposed, for which 33 parameters are estimated using full-information maximum-likelihood methods with filtering. Discrepancies in model performance are highlighted, which typify the difficulties of a filtering approach to water quality model development. The theoretical advantages of filtering, and its practical limitations, are summarized for water quality applications. It is concluded that filtering techniques offer valuable organizing concepts, and are themselves a valuable tool when dynamic quantification of the variance in state estimates is required. Filtering is most applicable when dealing with a low-dimensional system with a well-known model and dense data base, for which highly accurate short-term forecasts are required. However, less accurate models and estimation techniques may provide more cost-effective solutions to many water quality problems.

INTRODUCTION

Concern over the potential water quality effects of man's activities motivates increasingly rigorous management of water uses. As a result, demands are escalating for increasingly accurate water quality models, as management tools. The available data bases are becoming large, but are typically incomplete and contain measurement uncertainties. Use of this data for verification of dynamic, multivariable models is difficult using the more traditional qualitative and least-squares methods.

Estimation techniques using Kalman filtering and its extensions offer a potential solution to the problem of model development with noisy, incomplete data. In the past, these statistical methods have been refined and successfully applied to analogous problems in navigation and guidance, and industrial process control. These techniques have recently begun to be applied in water quality modeling. The purpose of this paper is to review the theoretical promise of these estimation techniques, tempered with a review and case study of practical experience gained in recent applications to water quality modeling.

THEORETICAL APPLICATIONS OF FILTERING

The Kalman filter is a statistical method for computing the best (i.e., minimum-uncertainty) state estimates for an uncertain linear system (17). Information about a system is acknowledged to exist both in a model of the system and in measurements of the system. It is also acknowledged that there are imperfections and/or stochastic effects in both the model and the data. Under these conditions, a best estimate of the system states, parameters, or other attributes is obtained by combining the information from the model and data. By formulating the system model in a stochastic form, model uncertainties can be quantified. Likewise, uncertainties due to noisy measurements can be quantified. When combined, the information from each of these sources is weighted according to its uncertainty; the filtered state estimate is the minimum-uncertainty "weighted-average" of model and measurement data. The filter also quantifies the variance in this state estimate.

Kalman filtering, with extensions, can be used for several estimation tasks: state estimation, parameter estimation, and model identification. (Hereafter, the group of filtering algorithms including the Kalman filter and extensions to non-linear systems will loosely be referred to as "filtering.") By analogy to earlier aerospace problems, several water quality applications are possible:

STATE ESTIMATION

Filtering offers very promising application to water quality monitoring. Utilizing any available measurements (no matter how scattered or asynchronous), the filter can optimally estimate water quality conditions throughout the modelled region. By "interpolating" between monitoring stations in this way, and by admitting irregular sampling, much more cost-effective monitoring may be possible. In addition, the filter can operate on-line, to process incoming measurements and update its water quality estimates in real time. This feature opens possibilities for real-time enforcement of quality standards and/or real-time control of water consumption and treatment.

PARAMETER ESTIMATION

Extensions of filtering theory offer powerful techniques for estimation of model parameters (8, 17). Methods are available to estimate model coefficients, initial conditions, boundary conditions, inputs, and the statistics of inherent model or measurement noise. Parameter estimation can be used for accurately quantifying one aspect of a modeled process (e.g., a nutrient uptake rate). Conversely, parameter estimation can "tune" an entire model for maximum accuracy in a specific application. Two parameter estimation techniques have developed from Kalman filtering: the Extended Kalman Filter and maximum-likelihood estimation (17).

MODEL IDENTIFICATION

Further extensions of filtering theory offer techniques for evaluating the underlying model structure for a system (8, 17). Both quantitative and qualitative measures of model performance are generated. These measures allow explicit, iterative identification of improved model structures, in an integrated model development process. These measures also provide useful *a priori* indications of a model's predictive capability.

EXPERIENCE WITH FILTERING IN WATER QUALITY MODELING

The preceding section shows that many water quality modeling problems may be posed as estimation problems which are theoretically solvable via filtering. This section discusses the practical limitations of filtering in water quality applications. Generic problems common to many environmental applications of filtering are discussed, along with potential remedies. Following this review, a recent application of filtering to hydrothermal model development is presented, as a case study in the advantages and disadvantages of these techniques.

GENERIC PROBLEMS

There is only limited experience with filtering in water quality applications. Most such references deal with proposed pollutant discharge control schemes involving filtered state estimates (1, 2, 4, 18). A few studies use filtering in parameter estimation and/or model identification (3, 5, 15, 18). There have also been similar applications in such related areas as groundwater flow (11), hydrology (6), and air quality (7). In a slightly different vein, several investigators have also used filtering in the design of monitoring programs (10). Based on these studies, environmental modeling problems appear to share several characteristics which typically limit the filter's usefulness:

- High dimensionality results from the spatially-distributed nature of environmental systems, and rapidly leads to uneconomic computational burdens in the filter's matrix operations.
- Uncertain noise statistics increase the uncertainty in filter results, perhaps to the point where less accurate (and less costly) estimation methods are equally useful. Parameter estimation techniques may improve our noise estimates but these methods aggravate the filter's already large computational burden.
- Limited measurements constrain the filter's estimation capabilities, because isolated or infrequent measurements produce only local or temporary reductions in uncertainty.
- Large model-structure uncertainties often exist, even after parameter estimation is completed (15), and are typically much greater than measurement errors. Thus, estimation accuracy is often severely limited by model accuracy, despite use of the most rigorous filtering algorithms.
- Inconvenient model transformation is often required, to convert common models into filter-compatible form.

These difficulties suggest that a complex model and statistically rigorous filtering algorithm may be warranted only where the most accurate state and covariance estimates are required. Where maximum accuracy is not required, simpler models and/or approximate estimation methods may prove to be much more cost-effective. For example, a simpler, lower-dimensional model offers greater economy and convenience. Model identification methods based on filtering now offer greater flexibility in the use of these models. In a review of water quality model development methods, Beck urges consideration of "black-box" model structures for many applications where the more intuitive "internally-descriptive" structures are commonly used. "The use of the black-box model is particularly advantageous when *a priori* information on the physical phenomena governing the system dynamics is minimal; in this case the black box is literally a fair reflection of our knowledge of the system and the model is a first attempt at elucidating any observed dynamical relationships. Alternatively, where a mechanistic, or internally descriptive, model is available, but its form is so complex that it requires the characterisation of too many parameters from an insufficient number of data, an input/output model can yield equally useful results in forecasting or control system synthesis applications (3)".

In general, simpler low-dimensional models appear to be most useful where concern focuses on conditions at only a few points, where a dense data base is available for iterative model development, and where short-term predictions are desired. Simple models are easily adapted to the state-space formulation required for optimal parameter estimation, and are easily modified. Internally-descriptive models require a difficult and non-intuitive transformation to state-space form, in order to undergo filtering. Consequently, any model changes must also undergo this laborious transformation. The internally-descriptive structure thus appears better suited for description of wide spatial variations in water quality, where an appropriate black-box model is difficult to deduce. Discriminating between model structures is ultimately a problem of finding the minimal realization to adequately represent a system. No solutions to this problem are yet available for environmental models.

Less rigorous estimation algorithms are an alternate response to filtering difficulties in water quality problems. McLaughlin suggests several possibilities: new formulations of the recursive filter equations, multi-level decomposition of the filtering problem, "sub-optimal" (approximate) filter formulation, and stochastic approximation (11). Each of these techniques is applicable only to certain types of problems; neither a general solution, nor general criteria for applicability, are available yet.

HYDROTHERMAL MODEL DEVELOPMENT USING FILTERING: A CASE STUDY

Background

(This work is summarized from earlier reports by Schrader and Moore (15, 16).) Salem Harbor Station is a 750 megawatt steam-electric generating station located in Salem, Massachusetts, on the Atlantic Coast 40 miles north of Boston. The plant employs a once-through condenser-cooling system, withdrawing seawater from Salem Harbor, and discharging the heated water back into the Harbor (see Figure 1).

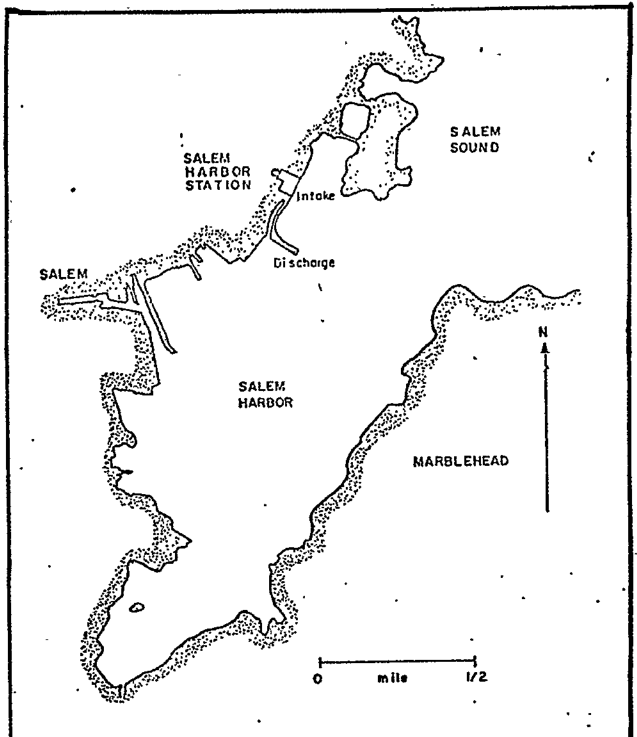

FIGURE 1 SALEM HARBOR STATION AND ADJACENT WATER BODIES

A discharge temperature limit established by pollution control authorities periodically forces uneconomic operation of the power plant. Consequently, a hydrothermal model of Salem Harbor is desired, to allow improved prediction and control of plant intake and discharge temperatures. The model development problem is posed as follows:

> Develop a simple and accurate model of intake temperatures at Salem Harbor Station. The model must predict up to 24 hours in the future. It must predict the daily peak intake temperature within +1^{o}F, 90% of the time. It must be developed using existing data from Salem Harbor.

To meet the model simplicity requirement, Salem Harbor is descretized as a two-basin, two-layer system (Figure 2). Heat and mass transfer models

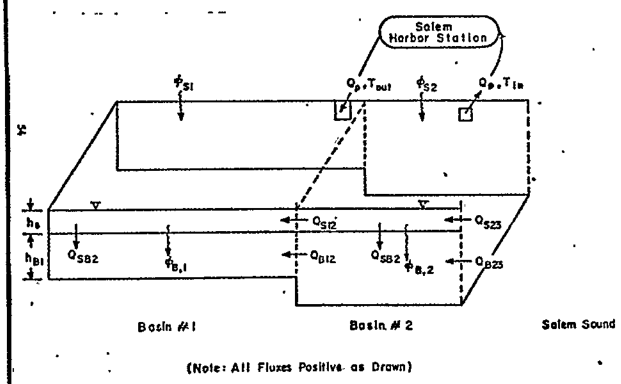

FIGURE 2 SCHEMATIZATION OF SALEM HARBOR FOR TWO-BASIN, TWO LAYER MODEL

are developed for this simplified representation. This "internally-descriptive" model (to use Beck's (3) terminology) is chosen in order to incorporate extensive previous hydrothermal modeling experience (9). (In retrospect, this formulation would not be chosen again for this type of problem, for reasons discussed in the following section.)

Limited data are available to develop and verify the model. Two years of nearly continuous hourly plant intake and discharge temperatures form the data core, with supporting hourly hydrological and meteorological data where available. Only limited temperature records exist for other points in the harbor; these are not included in this model development effort. Based on early estimates of computational cost, the data base for parameter estimation is confined to a 96-hour time series of intake temperature measurements. Four separate periods, comprising 106 days of additional data, are set aside for evaluation of the final model's performance.

Having selected a model structure and a data base, the model development problem posed above is reduced to a parameter estimation problem. Estimation is done using full-information maximum likelihood methods (14). This technique is chosen because it is theoretically optimal (given the "true" model structure), and likely to be less expensive than the extended Kalman filter for estimating the large number of unknown parameters involved in this problem. The technique is also very convenient to implement, using the GPSIE computer package (12, 13).

Within the chosen model structure, twenty-eight physical parameters are either unknown, poorly known, or intuitively assumed. This number of parameters arises because of the physical realism attached to individual coefficients, some of which could otherwise have been combined. In addition, five statistical parameters (describing model and measurement noise) are also unknown. Because of the stringent accuracy objective specified in the original model development problem, the global optimal estimate for all thirty-three parameters is sought.

The estimation process is divided into "preliminary" and "final" phases. In the preliminary phase, parameters are constrained to physically realistic limits, and subjectively varied to obtain the maximum-likelihood set of estimates. This manual search established initial estimates for the global optimization which are "near" the optimum. In the final phase, all twenty-eight physical parameters are estimated simultaneously, in two global estimations. A non-linear optimization routine (19) is used to iteratively converge on the maximum-likelihood parameter estimates. To complete the estimation, the five statistical parameters are subjectively re-estimated between the two estimations of physical parameters.

Results and Discussion

The estimation results are not what is expected. The model does perform qualitatively well. However, the model fails to meet the original accuracy criterion, and fails other statistical tests for acceptability. In addition, several parameters exhibit extreme and unrealistic estimates. These results, and their implications for future filtering applications, are discussed below.

Figure 3 presents an example of a 24-hour intake temperature forecast made with the final model. The cyclic rise and fall of predicted temperatures (due to the tidal cycle) agrees qualitatively with the amplitude and frequency of the observed temperature cycle. However, in this figure (and others) the model exhibits a chronic bias towards erroneously high temperatures, even after maximum-likelihood parameter estimation.

The complete sequence of all parameter estimates is too large to present here. Instead, Table 1 presents selected parameters, and their estimated values at various steps in the estimation process.

Parameters A1, A2 and hs exhibit physically unrealistic final estimates. A1 and A2 are the basin surface areas in the model (refer to Figure 3). Since the total Salem Harbor surface area is known to be 32-35 million square feet, (depending on tidal level), the final estimate of 12 million square feet is clearly unrealistic. Likewise, a surface layer thickness (hs) as large as 9 feet is not observed in Salem Harbor. Parameters g6 and F4 are other examples of extreme parameter estimates. In contrast, a few parameters (e.g. C2 and g3) do undergo the expected small perturbation from initial to final value. The model noise (σ_{S1}^2 and σ_{S2}^2) and measurement noise (σ_z^2) statistics in Table 1 characterize the noise levels existing during each reestimation of physical parameters.

Measures of model performance are computed for several test periods, using both the preliminary

FIGURE 3

PREDICTED INTAKE TEMPERATURES FROM FINAL MODEL 6/30 - 6/31/74

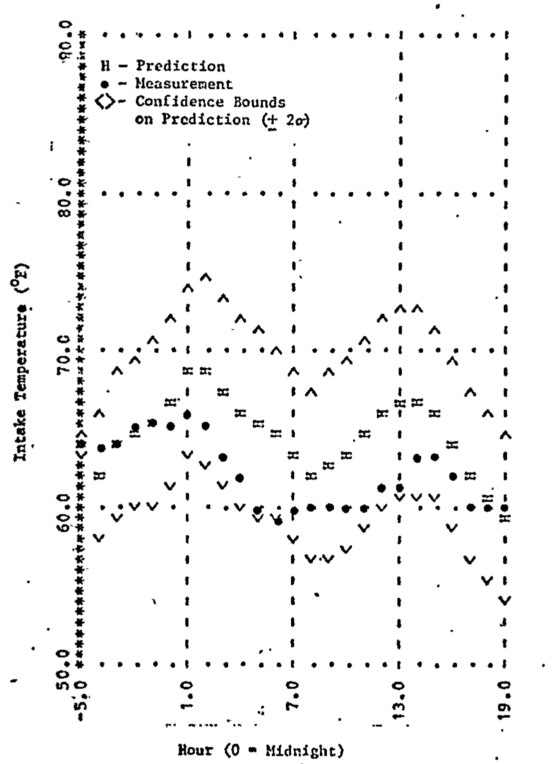

85

TABLE 1 - SELECTED PARAMETER ESTIMATES AT STEPS OF ESTIMATION PROCESS

PARAMETER (see text)	PRELIMINARY PHASE: INITIAL ESTIMATE	PRELIMINARY PHASE: BEST PHYSICALLY REALISTIC ESTIMATES	FINAL PHASE: FIRST GLOBAL ESTIMATES (PHYSICAL PARAMETERS)	FINAL PHASE: REVISED STATISTICAL PARAMETERS	FINAL PHASE: FINAL GLOBAL ESTIMATES (PHYSICAL PARAMETERS)
A1	20×10^6	20×10^6	19×10^6		4×10^6
A2	12×10^6	15×10^6	15.9×10^6		8×10^6
hs	2.	6.	6.82		9.83
C2	1.	1.	1.64		1.01
g3	.015	.015	.0157		.0145
g6	145.	133.	155.		-166.
F4	6.64	6.64	6.10		7160.
σ_{S1}^2	1.	5.		3.	
σ_{S2}^2	1.	20.		1.5	
σ_z^2	.04	.1		.04	

and final parameter values. Similar statistics are also computed for the base period over which the estimation was done. Selected results are presented in Table 2, for the following statistics:

- The log-likelihood is the statistic around which the parameter estimation is done, and thus should show the best improvement between preliminary and final models. A perfect model yields $\xi = 0$; progressively lower values reflect progressively worse models. Comparison is only valid between models run for the same time period.

T(n) - The percentage of daily peak intake temperature predictions within $n^{o}F$ of the measured peak. By definition of the model development problem, an acceptable model must yield T(1) 90%.

M - The mean-square error in 24-hour (unfiltered) temperature predictions.

P(a) - The auto-correlation ("whiteness") of measurement residuals, for lag a. A perfect model yields all P(a)=0; an acceptable model yields $P(a) \leq 5$ (14).

TABLE 2 - SELECTED MEASURES OF MODEL PERFORMANCE

	TEST PERIOD							
	6/25 - 7/22		7/29 - 8/21		8/23 - 9/20		BASE PERIOD 7/29 - 8/2	
	FINAL	PRELIM.	FINAL	PRELIM.	FINAL	PRELIM.	FINAL	PRELIM.
Log-likelihood ξ	-662.3	-691.9	-514.8	-568.4	-609.9	-650.3	-87.0	-105.9
T 1^{o}	16%	4%	5%	5%	27%	27%	--	--
2^{o}	24%	8%	28%	5%	54%	35%		
3^{o}	36%	20%	38%	19%	73%	54%		
M ($\sqrt{M}$)	19.4 (4.40)	38.6 (6.21)	16.4 (4.05)	26.9 (5.19)	10.1 (3.18)	7.7 (2.77)	10.7 (3.27)	27.3 (5.22)
Whiteness P(0)	- 3.9	-8.2	- 6.8	- 8.9	-8.7	-12.7	--	--
P(1)	12.3	8.9	7.1	5.4	7.1	2.8		
P(2)	5.8	5.5	2.8	3.7	2.8	0.9		
P(3)	0.6	3.7	0.3	2.4	0.2	0.4		

Based on Table 2, the model fails to solve the original model development problem, since T(i) is always less than 90%. The model also fails statistical tests on P(a), many of which are greater than five. Comparisons of preliminary and final model performance should show a consistent improvement in the final model, but do not. For example, the M statistic is better for the preliminary model during the last test period, and many of the P(a) statistics are better for the preliminary model. Furthermore, even where the final model performs better, the incremental improvement is small.

These results typify several generic difficulties with environmental applications of filtering: uncertain model structure and noise statistics, in construction with limited data; high cost; and inconvenience. Several results point to inaccurate model structure. The chronic bias towards erroneously high temperatures is visual evidence of structural flaws. Physically unrealistic values of certain parameters reflect the extremes to which the estimator must go, to compensate for structural flaws. Unacceptably high autocorrelation of residuals (P(a)) is another strong indicator of inadequate model structure. The small improvement between preliminary and final models also emphasizes that inadequacies in the final model are due to structural flaws, which parameter estimation cannot correct.

In this case, large uncertainties in the underlying model structure constrain the parameter estimation to a mediocre result. In retrospect, the modeling strategy may have focused too much on rigorous estimation of many parameters, and too little on iterative development of a simple but sufficient model structure. A lower-dimensional model would be consistent with the limited data base (one location) available for model development. A simpler, black-box model structure would also allow a more convenient model identification process; even the low-order internally-descriptive model used herein requires one week to effect minor changes. Because the covariances of model and measurement noise are not known, they must be estimated along with other parameters. However, maximum likelihood parameter estimation is sensitive to noise statistics (14); hence, the variation in parameter estimates shown in Table 1. During the second global estimation, noise covariances are much lower than during the first. In essence, this reduces the amount of model-data error which can be attributed to random variability, thus forcing the estimator to achieve a better "fit" of the data. It does this by forcing parameters to values more extreme than in earlier estimations. (It is likely that a broader or longer data base would reduce this effect; however, no test results are available.)

In addition to the statistical difficulties discussed above, full-information maximum likelihood parameter estimation is expensive for the moderate-sized internally-descriptive model used in this study. For example, one global estimation of twenty-eight parameters, over 96 data points, costs \$300 to \$500 on the MIT Information Processing Center's IBM 360/70. The filtering and maximum likelihood algorithms also have large core requirements, on the order of 300 k bytes.

SUMMARY AND CONCLUSION

Filtering techniques are theoretically applicable to several estimation problems in water quality modeling: state estimation (of water quality conditions); parameter estimation and model identification (for model development and verification). Experience with water quality applications is limited in both the number and size of problems studied. This experience indicates several difficulties common among environmental applications of filtering. High costs stem from the intricacies of filtering algorithms, the high dimensionality of spatially-distributed water quality systems, and the number of uncertain parameters requiring estimation. Uncertain noise statistics strongly influence parameter estimation and model identification. Scarce measurement data limits the filter's ability to improve parameter estimates and model structure. Large uncertainties in model structure may severely limit estimation accuracy.

In the future, the advantages and limitations of filtering should be evaluated on a project-by-project basis. However, the experience reported in this paper suggests the following guidelines:

1. Filtering techniques offer organizing concepts of major value for approaching the development and testing of water quality models in the presence of noisy, incomplete data.
2. Filtering techniques are necessary in applications requiring dynamic quantification of the variance in a system; the filter is the only available method for computing these values.
3. Filtering is best-suited to water quality modeling problems having low dimensionality, a dense data base, a well-known model structure, and for which highly-accurate short-term predictions are expected.
4. Approximate (and less accurate) estimation algorithms are available, which may provide a more cost-effective solution to many modeling problems, where: statistical rigor is not mandatory; the model or data is too weak to justify statistical rigor; or a high-dimensional system precludes statistical rigor for economic reasons.

ACKNOWLEDGEMENT

This research was supported by the New England Power Company, Westborough, Massachusetts.

BIBLIOGRAPHY

1. Beck, M. Dynamic Aspects of Water Quality Modeling and Control. CUED/F-CAMS/TR 130, Dept. of Engineering, University of Cambridge, England. (1976).
2. Beck, M. "Dynamic Modelling and Control Applications in Water Quality Maintenance". Water Research, Vol. 10. (1976).
3. Beck, M. Identification and Parameter Estimation of Biological Process Models. CUED/F-Control/TR 116, Department of Engineering, University of Cambridge, England. (1975).
4. Beck, M. "Problems of River Water Quality Modelling and Control: A Review", to appear in Proceedings, IFAC Symposium on Environmental Systems, Kyoto, Japan, August, 1977. (Pre-print - 1977).
5. Beck, B. and P. Young. "Systematic Identification of a DO/BOD Model Structure", Journal of the Environmental Engineering Division, ASCE, Vol. 102, No. EE5. (October, 1976).
6. Bras, R. L. and I. Rodriguez-Iturbe, Rainfall - Runoff as Spatial Stochastic Processes: Data Collection and Synthesis. Ralph M. Parsons Laboratory for Water Resources and Hydrodynamics, TR No. 196, MIT, Cambridge, Mass. (1975).
7. Desalu, A. A. Dynamic Air Quality Estimation in a Stochastic Dispersive Atmosphere, Electronic Systems Laboratory, Report No. ESL-R-542, MIT, Cambridge, Mass. (1974).
8. Eykhoff, P. (ed.), System Identification: Parameter and State Estimation, John Wiley and Sons, New York (1974).
9. Harleman, D. R. F. and K. D. Stolzenbach, Engineering and Environmental Aspects of Heat Disposal from Power Generation. Dept. of Civil Engineering, MIT, Cambridge, Mass. (June, 1975).
10. Lettenmaier, D. P. Design of Monitoring Systems for Detection of Trends in Stream Quality. Technical Report No. 39, Charles W. Harris Hydraulics Laboratory, Univ. of Washington, Seattle (August, 1975).
11. McLaughlin, D. B. "Application of Kalman Filtering to Groundwater Basin Modeling and Prediction." In Proceedings of the IIASA Workshop on Recent Developments in Real-Time Forecasting/Control of Water Resource Systems, International Institute for Applied Systems Analysis, Laxenburg, Austria (1977).
12. Peterson, D. W. and F. C. Schweppe, "Code for a General Purpose System Identifier and Evaluator (GPSIE)", IEEE Transactions on Automatic Control, Vol. AC-19, #6: 852-4 (Dec., 1974).
13. Peterson, D. W. GPSIE (General Purpose System Identifier and Evaluator) User's Manual. Pugh-Roberts Associates, Cambridge, Mass. (1976).
14. Peterson, D. W., Hypothesis, Estimation and Validation of Dynamic Social Models - Energy Demand Modeling, Ph. D. Thesis, Dept. of Electrical Engineering, MIT, Cambridge, Mass. (June, 1975).
15. Schrader, B. and S. Moore. Hydrothermal Modeling for Optimum Temperature Control: An Estimation-Theoretic Approach. Ralph M. Parsons Laboratory for Water Resources and Hydrodynamics, TR No. 214, MIT, Cambridge, Mass. (July, 1976).
16. Schrader, B. P., Once-Through Cooling at Salem Harbor Electric Generating Station: Short-term Temperature Forecasting, and Policy Aspects of a Temperature Standard, Environmental Engineer's Thesis, Department of Civil Engineering, MIT, Cambridge, Mass., (1976).
17. Schweppe, F. C., Uncertain Dynamic Systems. Prentice-Hall, Inc. Englewood Cliffs, N. J. (1973).
18. Young, P. and B. Beck. "The Modelling and Control of Water Quality in a River System", Automatica, (Sept., 1974).
19. Zangwill, W. I. "Minimizing a Function Without Calculating Derivatives", The Computer Journal, Vol. 10, No. 293. (1967).