Volume Title: Economic Analysis of Environmental Problems

Volume Author/Editor: Edwin S. Mills, ed.

Volume Publisher: NBER

Volume ISBN: 0-87014-267-4

Volume URL: http://www.nber.org/books/mill75-1

Publication Date: 1975

Chapter Title: Operational Problems in Large Scale Residuals Management Models

Chapter Author: Walter Spofford, Clifford Russell, Robert Kelly

Chapter URL: http://www.nber.org/chapters/c2836

Chapter pages in book: (p. 171 - 238)

Operational Problems in Large Scale Residuals Management Models

Walter O. Spofford, Jr., Resources for the Future, Inc.,
Clifford S. Russell, Resources for the Future, Inc., and
Robert A. Kelly, Resources for the Future, Inc.

Introduction

Over the past three years, we at Resources for the Future, Inc. have been working on the development of a regional residuals management model which in general form is in the classical mold, but which includes certain departures in detail that we consider important.[1] Like the classical models, it is designed to find the least-cost way of meeting ambient environmental quality standards given knowledge of the costs facing residuals dischargers and of the natural systems intervening between these dischargers and the points throughout the region at which quality is constrained. Unlike the earlier models, however, it is designed to deal with air and water quality and solid waste problems simultaneously because of the tradeoffs among airborne, waterborne and solid residuals implied by the conservation of mass and energy. In addition, we have developed industrial models, which are included as modules in the overall regional model, which

NOTE: We are grateful for the many helpful comments received from our colleagues Blair T. Bower and Allen V. Kneese, and from the conference discussant, J. Hayden Boyd.

1. A pathbreaking effort in this field was the work of the Delaware Estuary Comprehensive Study funded by the federal government to provide the basis for choosing stream standards and setting effluent standards (load allocations) in the Delaware Estuary. For a description of their model see Federal Water Pollution Control Administration, *Delaware Estuary Comprehensive Study* (Philadelphia, Pa.: U.S. Department of the Interior, July 1966).

can reflect the impact on residuals generation of changes in product mix, raw material quality, etc., and which include methods other than end-of-pipe treatment for altering residuals discharges.[2] Finally, the model and its method of optimum seeking are designed to be flexible with regard to the kinds of models of the natural environment which can be used. Thus, in particular, we do not limit ourselves to the linear transformation functions which traditionally have been used to connect discharges and ambient concentrations, but allow for inclusion of more complex formulations, including nonlinear simulation models.

Our approach thus far has been to construct small, "didactic" versions of this framework in order to test and develop our ideas without running up tremendous computer bills or getting buried in mountains of data. Two didactic applications have been constructed and are reported elsewhere.[3] In the first (see footnote 3, Russell and Spofford, 1972), appropriate demand functions and economic damage functions associated with ambient residuals concentrations at various locations throughout the region were assumed to exist, the environmental models—air dispersion and water quality—were assumed to be linear, and the objective function was one of net regional benefits. The institutional framework envisioned for this case was a regional management authority with powers to set effluent charges or standards.

In a follow-up, but still didactic application (see footnote 3, Russell, Spofford and Haefele, 1972), the model was expanded to provide information on the sociogeographic distribution of costs and benefits associated with meeting different levels of environmental quality. It was applied to an hypothetical region similar to the first one, and linear environmental models were again employed. The institutional framework envisioned for selecting levels of environmental quality, and for subsequent implementation of policy, was a legislative body.

Both these applications are reported elsewhere, hence, there is no need to go into further detail here. By way of introduction, though, we do show the overall model framework schematically in figure 1.

2. See C. S. Russell, "Models for the Investigation of Industrial Response to Residuals Management Action," *Swedish Journal of Economics* Vol. 73, No. 1 (1971): 134–156.

3. See C. S. Russell and W. O. Spofford, Jr., "A Quantitative Framework for Residuals Management Decisions," in *Environmental Quality Analysis: Theory and Method in the Social Sciences,* Kneese and Bower, eds. (Baltimore: Johns Hopkins Press, 1972). For a discussion of the framework as modified for use in a legislative setting, see C. S. Russell, W. O. Spofford, Jr., and E. T. Haefele, "Residuals Management in Metropolitan Areas" (paper delivered at the International Economics Association Conference on Urbanization and the Environment, 19–24 June, 1972, Copenhagen).

Figure 1

Schematic Diagram of the Regional Residuals Management Model

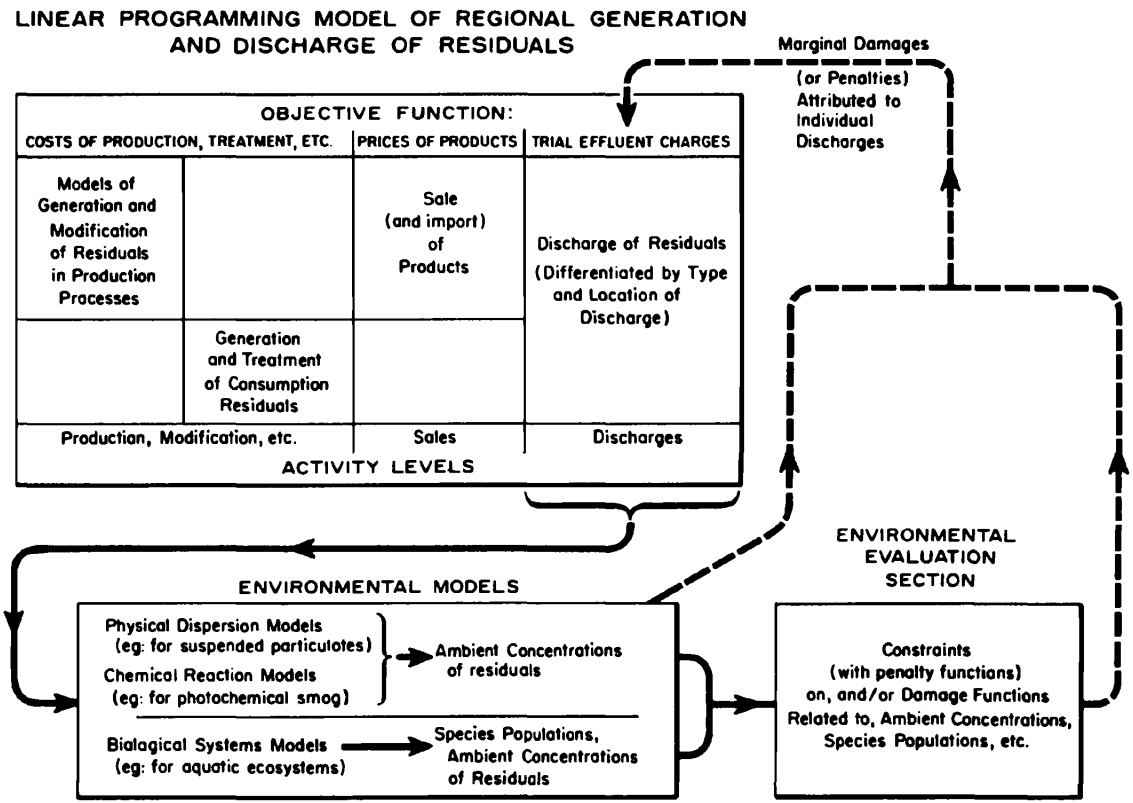

We have learned from our experience with didactic models that this approach is operationally feasible, at least for small scale applications. However, small scale applications to hypothetical regions provide us with very little indication of the operational difficulties involved in scaling up to an actual regional application in terms of the problems of collecting and subsequently manipulating massive quantities of data, and of the capability of present generation computers to cope with these large scale regional models. We are now at the stage of testing whether this framework can be applied to an actual region or whether it will become unmanageable when we attempt to deal with very large numbers of dischargers and locations throughout the region at which environmental quality is constrained. The question ultimately is whether we have developed a mildly interesting academic curiosity or a potentially useful management tool. To answer this question, we are now working on an application of the model based on the Delaware Valley region of New Jersey, Pennsylvania and Delaware, and this paper is a discussion of several of the important computational problems we are facing in this effort and of the different approaches we are exploring to overcome these problems. Some of these ideas and techniques are currently being tested in a first, relatively simple, version of the Delaware Valley application. We refer to this version as the Delaware Valley Base Model.

This particular model is deterministic and steady state, as were the two didactic versions. Only one season (which could represent either the "low flow" season or an entire year) is considered at a time. Also, from an economic point of view, the model is static. The main feature of this model is the inclusion of both nontreatment and on-site treatment management alternatives, along with nonlinear simulation models of the natural world, within an optimization framework. Options such as low flow augmentation, instream aeration, and regional sewage treatment facilities are not considered explicitly at this time. Later on we intend to expand upon this "base" model to include other management options which appear to be important but which have been neglected in this initial version of the Delaware Valley Model.

The optimum seeking technique that we are using is a form of the gradient method of nonlinear programming and involves iterating through a system of three submodels: (1) residuals generation and discharge submodels; (2) environmental submodels; and (3) an environmental evaluation submodel. This iterative process may be described, briefly, as follows. At iteration k, the generation and discharge submodel, which is structured as a linear programming problem, is solved using a set of effluent charges which is based on the state of the natural world

on the $(k-1)$th iteration. The resulting discharges are passed to the environmental models which transform them into information on ambient concentrations and species populations. These data on the resulting state of the natural world are then compared to exogenously specified standards of environmental quality. "Penalty" functions are used to reflect the solution's failure in meeting these standards; marginal penalties associated with each discharge of each type of residual are computed and returned to the generation and discharge model as prices on residuals discharges for the $(k+1)$st iteration. When all the constraints are met (within some predetermined tolerance) and no further improvement in the objective function is possible, successive sets of both discharges and effluent charges will be the same, and the algorithm has found an optimum.

Ultimately, we would hope that such a management model might be useful either to an executive agency, such as a regional environmental quality management authority, or to a legislative body. The model is purposely designed to be flexible enough to deal with environmental quality damage functions (if and when they are available) or sets of standards on ambient environmental quality. With nonlinear environmental models, meeting environmental quality standards is, as we shall see, more difficult computationally than employing economic damage functions. In this initial model, as a test for our optimization algorithm, we assume ambient standards must be met.

We shall report here on what we are learning from use of the Delaware Valley Base Model, and on some of the specific programming techniques we are using. It is hoped that these details will be of interest to others engaged in large-scale modeling projects.

Some Operational Problems of Large Scale Modeling Efforts

Models of residuals generation and discharge

Over the past decade, Resources for the Future, Inc. has conducted considerable research in the areas of industrial water use and residuals generation and discharge.[4] A number of linear programming models of industrial plants has been one of the outgrowths of this research program. These models include beet sugar plant, thermal electric plant, petroleum

4. See the paper in this volume by Blair T. Bower for a discussion of this research.

refinery, and integrated iron and steel production.[5] It has been our intention all along to include these models in the residuals generation and discharge portion of our Delaware Valley residuals management model. But the question of how best to do so has raised a number of practical problems. The major problem is model size as related both to round-off error in matrix inversion and to computer time required for solution. In this section, we shall discuss the pros and cons of two approaches for coping with the problem of size—decomposition, and construction of condensed models of the industrial plants.

Condensed Models of Industrial Plants. The full-scale industry models, which were developed for the individual industry studies, have the significant advantage of incorporating a large range of alternative responses open to the plant in the face of effluent charges or discharge standards. In addition, they make it possible to show how residuals generation and discharge, and response to management actions, change with such exogenous (to the regional residuals problem) influences as factor input costs, product mix, and available production—materials recovery —by-product technology. The problem is, of course, that the more the model incorporates, the larger it becomes. For example, the full-scale models developed for petroleum refining and steel production have between 300 and 500 rows. If we combined a number of these models into a single LP matrix by arraying the individual plant models along the diagonal, the resulting regional management model would exceed the computational reliability of the LP routines now available before even a fraction of a large, complex region's industries had been included. For the LP algorithm we are using (IBM's MPSX package), the upper limit on solution reliability is probably between 2,000 and 3,000 rows even though some have reported success with as many as 4,000 rows. As a general rule, though, for problems any larger than about 1,500 rows, care should be taken in checking and interpreting results.[6]

5. See Appendix II to *Future Water Demands: The Impacts on the Water Use Patterns of Selected Sectors of the United States Economy: 1970–1990,* a Study for the National Water Commission by C. W. Howe, C. S. Russell, and R. A. Young, assisted by W. J. Vaughan, all of Resources for the Future, Inc., June 1970; and *Residuals Management in Industry: A Case Study of Petroleum Refining,* C. S. Russell (Baltimore: Johns Hopkins Press, 1973).

6. These statements are based, in part, on the experiences of D. P. Loucks and D. H. Marks. There seems to be no agreement on the upper limit of the number of rows as it relates to solution reliability. Some have had trouble getting a reliable solution with as few as 1,500 rows. Others claim to have been successful with as many as 4,000 rows. The upper limit on row size depends, among others, on the condition of the matrix of coefficients which can differ tremendously among problems. The condition of a

One possible way around this problem is the construction of condensed, or collapsed, versions of the full-scale plant models. The condensation process consists of the following:

1. a choice of a limited number of important inputs, outputs (products), and residuals which would determine the number of rows in the new model;
2. a repeated solution of the larger, full-scale model for different residuals discharge constraint sets, as well as for different constraints on inputs or outputs;
3. a characterization of each solution as a vector with entries in the rows determined in 1; (these entries would be reduced in proportion to some standard unit of input or output, i.e., a natural unit for the petroleum refinery is a barrel of crude oil charged.)
4. an expression of the objective function value from the full-scale model's solution in terms of the same standard unit chosen in 3;
5. an addition to the set of summary vectors just derived, the necessary explicit discharge activities to which trial effluent charges may be attached.

Care must, however, be taken in the developmental stage of a condensed model to anticipate the subsequent price stimuli to be used in actual operation of the overall regional management model. The unit costs used as objective function entries for the summary vectors, and the additional stimuli to be applied in the regional model, are intimately related. Objective function entries for the summary vectors should comprise only those costs (and prices) which will not be accounted for explicitly when the condensed models are included as modules within the overall regional management model. Residuals discharges, for example, are priced separately in the regional model. Hence, in developing the condensed models, zero prices are used on these activities in the full-scale industry models. This insures that the objective function entries for the summary vectors do not include any charges for residuals discharges.

We have investigated this technique for making use of full-scale plant models; the report on the Delaware Valley Base Model detailed later in this paper includes collapsed models of two petroleum refineries. However, we have found problems with this approach. The most important one is that in order to duplicate even a fraction of the flexibility of the full-scale model, we must include a very large number of columns (i.e.,

---

matrix can usually be improved by proper scaling. For a more extensive discussion of this point, see W. Orchard-Hayes, *Advanced Linear-Programming Computing Techniques* (New York: McGraw-Hill, 1968), Chapter 6.

alternative solutions) in the condensed version. Considering, for a moment, only a single residual, the response of the condensed model to an effluent charge will more closely approximate the full-scale version the finer the grid of discharge constraints on which the condensation is based. But, it is important to note that this same statement applies to the multidimensional space containing the vector of all residuals of interest. If we have five residuals of interest, and if we confine ourselves to a very rough grid (e.g., high, medium and low levels of discharge), there are still 243 ($3^5 = 243$) alternative solutions of the full-scale model to be obtained and expressed in the appropriate vector form. If we increase the grid fineness to four levels of discharge, we increase the number of solutions and, hence, summary vectors to 1,024. It is clear that the expense and bookkeeping involved in constructing collapsed models is considerable and that their column size can become very large.

Whether or not the condensed model approach would be a solution to the row-size problem depends, of course, on the number of rows in the resulting condensed models, and on the number of significant residuals dischargers in the region. If the average size of the condensed models could be kept to ten rows, and if we ignore the requirements for artificial bounds for the step-size selection part of the overall solution method (to be discussed later), we could construct a single LP model for a region of between 150 and 200 dischargers. But keeping these condensed models to ten rows is not easy. Thus, if we wish to include only one input, one output (product), six primary residuals and two secondary residuals (sewage sludge and solids from particulate removal, for example), we are up to ten rows. Every refinement on the product or residuals side reduces the number of individual sources we can include. And we cannot, of course, neglect the necessity for artificial bounds, so that our "capacity" is very much lower than 150–200 plants; a guess would be 40–50. Now, for many regions this would be sufficient, but in a large industrialized region, such as the Delaware Valley, this would not begin to cover the significant sources of air and waterborne residuals, particularly when we realize that at least the largest municipal incinerators and sewage treatment plants have also to be included and provided with discharge reduction alternatives.[7]

Decomposition. Another alternative for dealing with model size is to

7. Metropolitan Philadelphia Interstate Air Quality Control Region, "Inventory of Emission Sources," Office of Air Programs, Environmental Protection Agency lists about 300 individual industrial plants, plus another 30 or so municipal incinerators and large institutional heating plants.

subdivide the large regional LP problem, which would be created by lumping all sources of residuals for which discharge reduction alternatives are available, into a series of smaller linear programs. This amounts to recasting the problem in a standard decomposed form:[8]

$$\min \quad \{c_1X_1 + \cdots + c_nX_n\}; \tag{1}$$

$$\text{s.t.} \quad A_{11}X_1 \qquad\qquad \geq b_1$$

$$\vdots \tag{2}$$

$$A_{n1}X_n \geq b_n$$

$$A_{12}X_1 + \cdots + A_{n2}X_n \leq b_{n+1}. \tag{3}$$

When there are no shared constraints (equation 3) among individual "decomposed" components, each of the smaller LP's may be solved, in turn, as a separate subproblem prior to entering the environmental model subroutines for determination of resulting ambient environmental quality. This is, in fact, the case with the Delaware Valley Base Model which was purposely divided into two LP's to be solved separately. Dividing the optimization problem up this way is conceptually straightforward and certainly appealing from a computational point of view, even though there are certain practical difficulties in using the MPSX routine in this manner. But these difficulties are primarily matters of keeping MPSX outputs and inputs straight when there are many discharges, artificial bounds (step sizes), and trial effluent charges to be passed back and forth among various LP's and FORTRAN subroutines.

In a more complex model of a region, it may be impossible to provide individual, unconnected LP subproblems. This, of course, depends upon the interconnections—both market and nonmarket—among activities in the region. For example, if the environmental models were linear, and if they were dealt with in the regional model as part of the constraint set, it would be virtually impossible to subdivide the regional model into separate, unconnected submodels. However, even the elimination of environmental models as part of the constraint set does not guarantee us that we will be able to subdivide the regional model into a series of sepa-

8. For a discussion of the decomposition principle for linear programming, see G. B. Dantzig, *Linear Programming and Extensions* (Princeton: Princeton University Press, 1963) or G. H. Hadley, *Linear Programming* (Reading, Mass.: Addison-Wesley Publishing Co., Inc., 1962).

rate submodels. There are other types of relationships that inherently link activities together. For example, in our simple models, such as the Delaware Valley Base Model, we have had a market link between the petroleum refineries and home heating through the purchase of various grades of distillate fuel oil. The form of these constraints has been a simple one: production plus imports must be greater than or equal to regional use. (Heating for domestic purposes has been assumed price inelastic.)

The obvious approach to this problem is to treat the set of linear models with shared linear constraints as a classical decomposed linear programming problem and to solve it as such before entering the environmental models. Since decomposition algorithms are themselves iterative, we would be building in a set of iterations within each iteration of the overall management model, and this may involve us in a significant increase in computation time. A major drawback of this approach is that the LP algorithm does not have a decomposition algorithm built into it. Hence, to take advantage of decomposition, we would have to improvise. We are presently exploring this possibility.

In summary, the main concern we have with models of regional generation and discharge of residuals, when included as part of a regional residuals management model of a large, complex region, is with sheer model size. Our concern relates not only to the problems associated with round-off errors, but also to computational time and expense. We have proposed various approaches to the size problem and have investigated many of them using the Delaware Valley Base Model. Currently, it appears that a combination of collapsed versions of the full-scale industry models, standard decomposition, and sequential solution of a set of LP submodels is a feasible approach to the problem of attaining solutions to large scale, regional residuals management models.[9]

9. There is, at least in principle, a third possibility for reducing model size. All of the industry models which we intend to include in our regional residuals management model are linear, but their constraint sets contain a significant number of equality constraints. Each equality constraint could be used to eliminate a variable (column) in the LP. But once an industry model is built, this is a time consuming procedure and errors are likely to result. In addition, some of the eliminated variables are likely to provide useful information for management decisions and, hence, would have to be computed anyway after the optimization phase of the analysis were complete. Thus, the choice of variables to be retained is most important. For example, it would not be desirable to eliminate the residuals discharge vectors because both discharges (which are input to the environmental models) and prices on these discharges change from iteration to iteration. Although we recognize elimination of equality constraints as a possibility for dealing with model size, up to this point in our research, we have not given it serious consideration. In the future, however, we may.

Environmental models

Environmental models—air and water dispersion, chemical reaction, and biological systems—are used to describe the impact on the natural environment of energy and material residuals discharged from the production and consumption activities of man. We use these models to predict steady state concentrations of residuals and related substances (e.g. algae, oxygen in the estuary) at various points in the regional environment, given: (a) a set of residuals discharge levels from the linear programming submodel of regional generation and discharge of residuals; and (b) a set of values for the environmental parameters such as stream flow and velocity, wind speed and direction, atmospheric stability, and atmospheric mixing depth.

Some environmental models are easier to deal with than others within an optimization framework. It depends, in general, upon the mathematical structure of the model. In terms of the complexity involved, we find it useful to distinguish among four broad categories: (1) linear, explicit functions; (2) linear, implicit functions; (3) nonlinear, explicit functions; and (4) nonlinear, implicit functions.

We are currently using two environmental submodels in conjunction with our regional residuals management model. The first, a linear atmospheric dispersion model, is used to predict ambient concentration levels throughout the region of sulfur dioxide and airborne particulates. It was provided to us by the Environmental Protection Agency.[10] The second, a nonlinear aquatic ecosystem model, used to predict various ambient concentrations in the estuary, was developed at Resources for the Future specifically for our Delaware Valley residuals management study.[11] The inclusion of these environmental models has hopefully increased the usefulness of the overall management model for purposes of better informing public policy, but has raised several computational problems also. These two environmental models represent the extremes of complexity for inclusion within a management framework. A discussion of each model will raise some of the important issues and will reveal some of the problems involved.

Atmospheric Dispersion Model. Of the various atmospheric quality models which are available now, physical dispersion models are the most advanced. Chemical reaction models, such as for photochemical smog, are

10. Division of Applied Technology, Office of Air Programs, Environmental Protection Agency, Durham, N.C.

11. R. A. Kelly, "Conceptual Ecological Model of the Delaware Estuary," *Systems Analysis and Simulation in Ecology* Volume IV (New York: Academic Press, 1974).

being developed, and carbon monoxide models for urban areas are appearing. The most successful modeling efforts to date have been associated with predicting both steady and nonsteady state concentration distributions of sulfur dioxide and suspended particulates of 20 microns or less in diameter. Because of the availability of an existing air dispersion model, we selected ambient levels of sulfur dioxide and suspended particulates to represent the air quality of our region.[12]

The atmospheric model which we are using is the air dispersion model from the federal government's Air Quality Implementation Planning Program (IPP).[13] This model uses a dispersion model developed by Martin and Tikvart which evaluates concentrations downwind from a set of point and area sources on the basis of the Pasquill point source, Gaussian plume formulation. The Gaussian plume formulation may be used to estimate ambient concentrations under deterministic, steady state conditions. For any given source-receptor pair, production process and abatement device, specified meteorologic conditions, and discharge rate of unity, this nonlinear equation reduces to a linear coefficient relating ambient concentrations with residuals discharge rates.

The necessary inputs to this model are: x-y coordinates of all sources and receptors in the region; emission rates for each source—point and area; physical stack height, stack diameter, stack exit temperature, and stack exit velocity for each point source; a seasonal joint probability distribution for wind speed, wind direction, and atmospheric stability; a mean seasonal temperature and pressure; and a mean atmospheric mixing depth for the period of interest.

The output of this air dispersion model represents arithmetic mean seasonal concentrations of sulfur dioxide and airborne particulates based on the probabilities of occurrence of 480 discrete meteorological situations. For this computation, 16 wind directions, 6 wind speed classes, and 5 atmospheric stability classes are considered for each source-receptor pair with the occurrence of all combinations possible (hence, $16 \times 5 \times 6 = 480$ total possibilities). The joint probabilities of occurrence for each of these 480 combinations are determined from actual meteorological data.

12. We should point out that the selection of sulfur dioxide and suspended particulates as measures of air quality in our model coincides with real world considerations. These are, in fact, the first two airborne residuals for which environmental quality standards have been set in the United States.

13. See TRW, Inc., *Air Quality Implementation Planning Program* Vols. I and II (Washington, D.C.: Environmental Protection Agency, 1970), also available from National Technical Information Service, Springfield, Virginia, 22151, accession numbers PB 198 299 and PB 198 300 respectively.

For a given set of meteorological conditions and physical parameters, the vector of mean seasonal concentrations of sulfur dioxide and particulates, R, may be expressed linearly, in matrix notation, as;

$$R = AX + B, \tag{4}$$

where X is a vector of sulfur dioxide and particulate discharge rates; A is a matrix of transfer coefficients which specify, for each source-receptor pair in the region, the contribution to ambient concentrations associated with a residuals discharge rate of unity; and B is a vector of background concentration levels. The matrix of transfer coefficients, A, is the output of the dispersion model.

The important thing to note from equation 4 is that the state of the natural world (R) is expressible directly in terms of linear, explicit algebraic functions. This particular mathematical form is relatively easy to deal with in an optimization framework. In fact, equation 4 in its present form may be incorporated directly within the constraint set of a standard linear program when one of the management objectives is to constrain ambient concentrations of residuals.

As we shall see in the next section, one of the requirements of our optimization scheme is the availability of an environmental response matrix, $\partial R_i/\partial x_j$, $i = 1, \ldots, m$; $j = 1, \ldots, n$; where m is the total number of environmental quality indicators at all the designated receptor locations in the region and n is the total number of residuals discharges in the region. This matrix may be obtained by differentiating equation 4 with respect to all the residuals discharges in the region. That is,

$$\left(\frac{\partial R}{\partial X}\right) = A. \tag{5}$$

Before we leave this section, we should point out that not all atmospheric quality models are as easy to deal with as the physical dispersion models which are expressed in linear, explicit analytical form. Chemical reaction models, such as for photochemical smog, for example, would be significantly more difficult to handle within our optimization framework. The kinds of problems we would face with them are revealed in the discussion of a nonlinear aquatic ecosystem model which follows.

Aquatic Ecosystem Model. There are a variety of indicators which are commonly used for describing the quality of a body of water. Among them are pathogenic bacterial counts (or counts of an indicator thereof),

algal densities, taste, odor, color, pH, turbidity, suspended and dissolved solids, dissolved oxygen, temperature, and population sizes of certain plant and animal species. Because of the importance of dissolved oxygen to virtually all species of higher animals, and the relative ease with which it can be measured and modeled for a river or estuary, its concentration has been, and still is, one of the most frequently used criteria for setting general water quality standards.

Streeter-Phelps type dissolved oxygen models have been used for many years to predict water quality as a result of discharges of organic material (most notably, sanitary sewage).[14] Given certain assumptions about the natural environment, these DO models can be expressed as a set of linear algebraic relationships analogous to the linear air dispersion models discussed previously. From a computational point of view, they are very easy to deal with. This is, in fact, one reason for their continued popularity.[15]

However, these models have three deficiencies which we feel warrant the exploration of more sophisticated aquatic ecosystem models. First, we are really interested in the dissolved oxygen level only insofar as it is an accurate indicator of such things as algal densities and the population sizes of certain species of fish. To the extent that these densities and populations can vary independently of dissolved oxygen concentrations, we need information about them if policies on water quality are to be established intelligently. Second, materials other than organics (for example, nutrients and toxics) are known to have significant effects on aquatic ecosystems. Consequently, these inputs should be included along with the organics in order to evaluate more fully the impact on the environment of residuals discharges. Finally, systems ecologists feel that aquatic ecosystem models based on at least some biological (or ecological) theory, which includes the mechanisms of feeding, growth, predation, excretion, death, and so on, are more reliable for predicting dissolved oxygen levels than the more empirically based models of the Streeter-Phelps variety.

The aquatic ecosystem model we have developed is based on a trophic level approach.[16] The components of the ecosystem are grouped in classes

14. H. W. Streeter and E. B. Phelps, "A Study of the Pollution and Natural Purification of the Ohio River," *Public Health Bulletin No. 146* (Washington, D.C.: U.S. Public Health Service, 1925).

15. Models of the BOD-DO type are in widespread use. A typical example is given for the Delaware Estuary by R. V. Thomann in *Systems Analysis and Water Quality Management* (New York: Environmental Science Services Division of Environmental Research and Applications, Inc., 1972), pp. 160–81.

16. For examples of this approach, see R. B. Williams, "Computer Simulation of Energy Flow in Cedar Bog Lake, Minnesota, based on the classical studies of Linde-

("compartments") according to their function, and each class is represented in the model by an endogenous, or state, variable. Eleven compartments are designated in our model. The endogenous variables representing these eleven compartments are nitrogen, phosphorus, turbidity (suspended solids), organic material, algae, bacteria, fish, zooplankton, dissolved oxygen, toxics, and heat (temperature). In addition, the following exogenous variables (parameters) are considered: turnover rate (or advective estuary flow), and inputs (of the eleven chemical and biological materials above). Carbon is assumed not to be limiting and, hence, is not considered as either an endogenous or exogenous variable. Material flows among compartments within a given reach of the river (or estuary) are depicted in figure 2.

Figure 2

Diagram of Materials Flows Among Compartments Within a Single Reach

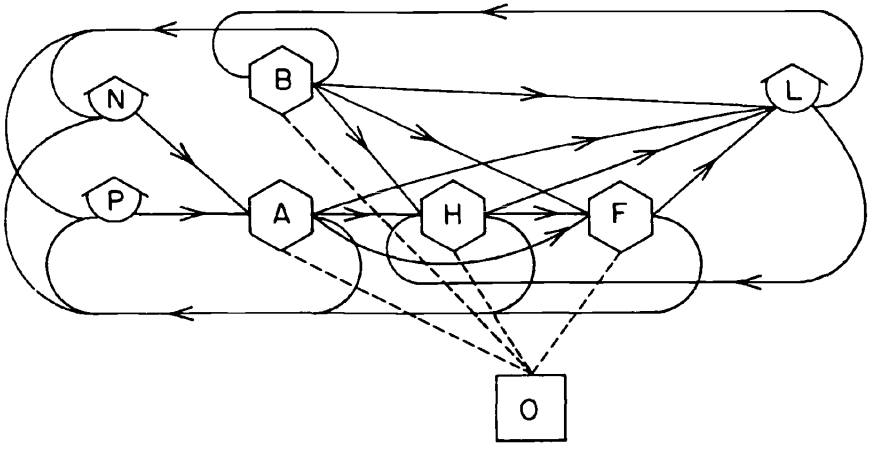

The inputs to the estuary model from the residuals discharges in the region are organic material measured by its BOD, total nitrogen, phosphorus, phenols (toxics), and heat. The outputs of this model of concern

---

man," *Systems Analysis and Simulation in Ecology* Vol. I, B. C. Patten, ed. (New York: Academic Press, 1971), pp. 543–82, and H. T. Odum, *Environment, Power and Society* (New York: Wiley-Interscience, 1971).

to us are densities of fish biomass, algal densities, and dissolved oxygen levels. The levels of these outputs are constrained; that is, environmental standards are imposed. In addition, concentrations of nitrogen, phosphorus, suspended solids, and organic material; temperature; and mass of bacteria and zooplankton are also available as by-product outputs of this model.

The time rate of change of material in each compartment is expressed in terms of the sum of the transfers among other compartments, and between adjacent sections of the estuary (since the material is distributed spatially as well as temporally). To insure mass continuity of the materials considered, material entering and leaving a compartment is explicitly accounted for.[17] The mathematical description of material transfers among compartments is based on the theoretical-empirical formulations given by Odum.[18]

Each compartment requires a separate differential equation to describe mass continuity, and in general, these equations must be solved simultaneously. In this particular case, the differential equations are ordinary ones of the first order, nonlinear variety. A set of similar differential equations is required for each reach of the estuary.[19]

The general form of the differential equation set for the kth reach may be expressed as,

$$\left(\frac{dR}{dt}\right)^k = f[R(t)^{k-1},\ R(t)^k,\ R(t)^{k+1},\ X(t)^k], \tag{6}$$

where $(dR/dt)^k$ is a vector of time rates of change of the endogenous variables R^k in the kth reach, and X^k is a vector of residuals discharges into the kth reach.

17. For the two nutrients—phosphorus and nitrogen—a mass balance is made on the individual chemical elements. For species, a mass balance is made on the total weight of carbon, nitrogen and phosphorus (assuming a constant ratio among them, i.e., C:N:P = 40:5:1 for all species).

18. H. T. Odum, "An energy circuit language for ecological and social systems: its physical basis," in *Systems Analysis and Simulation in Ecology* Vol. II, B. C. Patten, ed. (New York: Academic Press, 1972), pp. 139–211.

19. Finite difference forms of the more general partial differential equation set for describing mass continuity are used for the distance (space) variable. This is why we are able to write a separate set of differential equations for each reach (section) of the estuary. However, within each reach, time is expressed continuously (thus, the set of total differential equations rather than algebraic equations). When these differential equations are solved with analog computers, the concentrations of materials are continuous in time. When these equations are solved using digital computers, time must also be expressed in the finite difference form, and in the process, they reduce to a simultaneous set of nonlinear algebraic equations.

There are two problems associated with this ecosystem model formulation which we wish to discuss in more detail: (1) that of obtaining a steady state solution, and (2) that of obtaining an environmental response matrix. The first relates to models of this type in general whereas the second relates only to those situations where ecosystem models are to be included within an optimization framework.

Solution Methods. In its present form, equation 6 represents a set of ordinary nonlinear differential equations—one equation for each compartment, and one set of compartmental equations for each estuary reach —which must be solved simultaneously. If we were interested in the transient (or nonsteady) states of the system, simulation techniques, i.e., numerical integration (simulating first over space and then time) provide us with a readily available means of solution. However, we are interested only in the steady state solution.

For determining steady state solutions, there are two possibilities (or a combination thereof), neither of which guarantees finding a stable point equilibrium: (1) simultaneous simulation of a nonlinear differential equation set, and (2) simultaneous solution of a set of nonlinear algebraic equations. If we neglect inputs to, and outflows from, each reach due to longitudinal dispersion, the system can be dealt with first over time, and then space, starting with the uppermost reach and progressing systematically down the estuary.[20] In this case, equation 6 for the kth reach would reduce to,

$$\left(\frac{dR}{dt}\right)^k = f[R(t)^{k-1},\ R(t)^k,\ X(t)^k]. \tag{7}$$

Now, only the eleven compartmental equations within each reach must be solved simultaneously. The state of the system within a particular reach depends only upon the inputs from upstream, $R(t)^{k-1}$ and the residuals discharges to the kth reach, $X(t)^k$, both of which may now be treated as exogenous inputs. In addition, if the resulting steady state

20. Neglecting longitudinal dispersion, even in an estuary, is not as unreasonable as it first appears. Finite difference techniques for solving these differential equations introduce a numerical diffusion effect into the model. Inputs are immediately mixed in the volume, not because of any physical effects, but solely because of the numerical procedure. See D. J. O'Connor and R. V. Thomann, "Water Quality Models: Chemical, Physical, and Biological Constituents," *Estuarine Modeling: An Assessment* (Washington, D.C.: Environmental Protection Agency Stock No. 5501-0129, February 1971), Chapter III, p. 138.

solution, $(R^*)^k$, is independent of the time paths of rates of inputs, $R(t)^{k-1}$ and $X(t)^k$, equation 7 reduces to,

$$\left(\frac{dR}{dt}\right)^k = f[(R^*)^{k-1}, (X)^k, R(t)^k]. \tag{8}$$

Usually, ecological models are solved by simulation. Simulation of the differential equation set (a set of equations similar to equations 6 through 8 poses no particular problem, but the steady state solution, if one exists at all, may take considerable time. Oscillations can, and do, occur, and solutions may be otherwise unstable; they may become infinitely large. However, May[21] has demonstrated for a set of reasonable assumptions and a similar predator-prey nonlinear model, that these systems possess either a stable point equilibrium or a stable limit cycle.

Even when a steady state solution can be found, an additional problem is that there may be more than one stable point equilibrium. To investigate this problem, we ran an experiment with our ecosystem model. We used a random number generator to provide us with a set of random starting points. Twenty-five random starts resulted in the same steady state solution which indicates that our model is probably well behaved in this respect. However, another model may not be.

At steady state, $dR/dt = 0$, and thus the differential equation set above, equation 8, reduces to a set of nonlinear algebraic equations of the following form.

$$0 = f[(R^*)^{k-1}, (X)^k, (R^*)^k]. \tag{9}$$

The endogenous variables, $(R^*)^k$, are implicitly expressed in this formulation.

Various numerical methods, such as Gauss-Seidel and Newton's, have been used with success for solving simultaneous nonlinear algebraic equation sets, but each has its faults. Gauss-Seidel (also known as "Successive Approximation") has slow convergence properties, but it is relatively stable. Newton's method has more rapid convergence properties, but it is sensitive to initial conditions and it is often unstable.

Determination of steady state values for the endogenous variables in a nonlinear ecosystem model is difficult due to the nonuniqueness and

21. R. M. May, "Limit Cycles in Predator-Prey Communities," *Science* Vol. 177 (September 1972), pp. 900–902.

complexity of the solution. Even the stability characteristics of the steady state solution cannot be determined prior to its solution. With linear models, we can solve for the eigen (or characteristic) values of the differential equation set. These will tell us whether or not the time independent solution converges to a finite set of values, or diverges to infinity, or even if oscillations are involved—stable, diverging, or converging. For the nonlinear differential equation set, the best we can do is linearize the system at some point, and examine the eigen values of the resulting linear form. But this only tells us what is happening locally.

At this time, we are using a combination of Newton's method and simulation. These techniques are being used in the following way. Starting with the first reach, a solution is attempted by Newton's method using an estimate of the steady state values of the endogenous variables as an initial point. If a steady state solution is obtained, a solution for the second reach is attempted, using the steady state values of the endogenous variables from the first reach as a starting point. This procedure is repeated until the steady state solution is obtained for the last reach, or a reach is encountered which cannot be solved by Newton's method. When a solution cannot be obtained, an approximate solution is generated by numerically integrating the equation set over a one hundred day period. Empirical observation of the solution behavior indicates this is a fairly decent steady state solution. The simulation solution is then used to solve the next reach by Newton's method, and so on.

The Ecosystem Response Matrix. To include this nonlinear aquatic ecosystem model within the residuals management model, in addition to determining a set of steady-state values, it is also necessary to evaluate the response throughout the ecosystem to changes in the rates of the residuals discharges. That is, it is necessary to know, for example, the effect on algae in reach 17 of an additional BOD load discharged into reach 8, and so on. This requirement results in a considerable number of additional computations, but this knowledge of the system response, in conjunction with the penalty functions to be discussed in the next section, is the key to being able to use these complex ecosystem models within the optimization framework.

The response matrix we wish to compute may be expressed in matrix notation as, $\partial R/\partial X$, where R is a vector describing the state of the system throughout the entire length of the estuary, and X is a vector of residuals discharges throughout the region. Using equation 9 for each reach of the estuary, and the relationship,

$$Z^k = q^k R^{k-1} + X^k, \tag{10}$$

where Z^k is a vector of inputs to the kth reach, R^{k-1} is a vector of concentrations of materials in reach $k-1$, q^k is the estuary advective flow rate into the kth reach, and X^k is a vector of residuals discharges to the kth reach, a section of the system response matrix may be computed accordingly,

$$\frac{\partial R^i}{\partial X^i} = \frac{\partial R^i}{\partial Z^i} \cdot \frac{\partial Z^i}{\partial R^{i-1}} \cdot \frac{\partial R^{i-1}}{\partial Z^{i-1}} \cdots \frac{\partial R^{j+1}}{\partial Z^{j+1}} \cdot \frac{\partial Z^{j+1}}{\partial R^j} \cdot \frac{\partial R^j}{\partial Z^j} \cdot \frac{\partial Z^j}{\partial X^j}. \tag{11}$$

From equation 10 we note that,

$$\frac{\partial Z^k}{\partial R^{k-1}} = q^k I, \tag{12}$$

and,

$$\frac{\partial Z^k}{\partial X^k} = I, \tag{13}$$

where I is the identity matrix. Thus, the $(Z^k/\partial R^{k-1})$ terms are known a priori and are exogenous parameters in the ecosystem model.

The other terms, $(\partial R^k/\partial Z^k)$, are evaluated from equation 9 according to the rules for differentiating implicit functions.[22] That is,

$$\frac{\partial f}{\partial R} \cdot \frac{\partial R}{\partial Z} = -\frac{\partial f}{\partial Z}, \tag{14}$$

or,

$$\frac{\partial R}{\partial Z} = -\left(\frac{\partial f}{\partial R}\right)^{-1} \frac{\partial f}{\partial Z}. \tag{15}$$

This operation involves the inversion of the Jacobian matrix $(\partial f/\partial R)$. In addition, because the system of equations is nonlinear, the Jacobian matrix $(\partial f/\partial R)$ must be recomputed for each resulting state of the natural world.

It should be clear, then, from the above discussion, that the major problem associated with including environmental models within our management framework is one of computer time. Nonlinear representations of the natural world increase the complexity and the number of

22. See, for example, I. S. Sokolnikoff and R. M. Redheffer, *Mathematics of Physics and Modern Engineering* (New York: McGraw-Hill Book Co., Inc., 1958), pp. 237–41.

calculations necessary for each iteration, but hoping that they will also increase both the realism and predictive capability of the model.

Management model formulation and optimization scheme

In this section, we (a) present a formal mathematical description of our regional residuals management model; (b) indicate the method of handling certain kinds of constraints which are difficult, in fact in some cases impossible, to deal with in the traditional manner; and finally (c) we discuss the optimization procedure we are using.

Model Formulation. The objective function we are currently using is expressed in the form of a net benefit function. Hence, the objective is to maximize. The positive elements in this function include gross revenues from the sale of various products. The negative elements include: all the opportunity costs of traditional production inputs; all liquid and gaseous residuals modification (treatment) costs; and all collection, transport, and landfill costs associated with the disposal of solid residuals.

There are, basically, three types of constraints in the management model: traditional resource availability (inequality) constraints; continuity relationships (equality constraints); and residuals management (inequality) constraints. The latter, which involve the use of environmental models, are employed to constrain the levels of ambient environmental quality. The nature of all three types of constraints has been discussed in detail elsewhere.[23] We do not elaborate again on the first two types here. The third type is discussed in a slightly different context, as we are now treating these constraints a little differently than we did before.

Before proceeding, let us state the residuals management problem formally.[24]

$$\max \{F = f(X\ R)\}; \tag{16}$$

$$\text{s.t.} \quad g_i(X) = 0 \qquad i = 1, \ldots, m < n - q, \tag{17}$$

$$g_i(X) \geq 0 \qquad i = m + 1, \ldots, p, \tag{18}$$

23. C. S. Russell and W. O. Spofford, Jr., "A Quantitative Framework," and C. S. Russell, W. O. Spofford, Jr., and E. T. Haefele, "Residuals Management."

24. Note that the environmental relationships could have been written directly as,

$$h_i(X) \leq S_i, i = 1, \ldots, q.$$

However, we choose to deal explicitly with the variables R_i, $i = 1, \ldots, q$, here as they will be useful to us in a later development.

$$h_i(X) = R_i \qquad i = 1, \ldots, q, \tag{19}$$

$$R_i \leq S_i \qquad i = 1, \ldots, q, \tag{20}$$

$$R_i \geq 0 \qquad i = 1, \ldots, q, \tag{21}$$

$$X_i \geq 0 \qquad i = 1, \ldots, n, \tag{22}$$

where $f(X\ R)$ is, in general, a nonlinear objective function; $g_i(X) = 0$, $i = 1, \ldots, m$, is a set of linear equality constraints; $g_i(X) \geq 0$, $i = m + 1, \ldots, p$, is a set of linear inequality constraints; $h_i(X) = R_i$, $i = 1, \ldots, q$, represents a set of environmental functions which relate ambient concentrations of residuals to residuals discharges; X_i, $i = 1, \ldots, n$, is a vector of decision variables, including residuals discharges; R_i, $i = 1, \ldots, q$, is a vector of ambient levels of residuals concentrations and population sizes of species; and S_i, $i = 1, \ldots, q$, is a vector of ambient environmental quality standards (e.g., sulfur dioxide and particulates in the atmosphere, and algae, fish, and dissolved oxygen in the water).

As we have pointed out previously, some of the necessary environmental functions $h_i(X) = R_i$ are available in linear form (e.g., the air dispersion relationships and the Streeter-Phelps type dissolved oxygen models).[25] Others are only available in nonlinear analytical form, while still others are available in various other forms. As we pointed out in our discussion of nonlinear aquatic ecosystem models, no analytical expressions for them —either linear or nonlinear—of the form $h(X) = R$ are available. The variables R_i, $i = 1, \ldots, q$ are expressible only as a set of implicit nonlinear functions and, hence, simulation and other iterative techniques must be used to compute their values. From this discussion, we note that, in general, the environmental constraint set, equation 19, represents a variety of functional forms, many of which are difficult, or even impossible, to deal with using traditional mathematical programming techniques.

Because our optimization scheme, to be described below, requires that all the constraints be linear, we remove the environmental relationships from the constraint set and deal with them in the objective function. This modification of the problem requires the use of the penalty function concept which we shall discuss below.[26]

25. When environmental functions are expressible in this particular linear analytical form, their coefficients are known in the literature as transfer coefficients.

26. The use of "penalty functions" for eliminating constraints is not a new idea. It is a well-known technique and is in frequent use in one form or another under a variety

The new optimization problem may be stated formally as,

$$\max \{F = f(X) - P(X)\}; \tag{23}$$

$$\text{s.t.} \quad g_i(X) = 0 \qquad i = 1, \ldots, m, \tag{17}$$

$$g_i(X) \geq 0 \qquad i = m + 1, \ldots, p, \tag{18}$$

$$x_i \geq 0 \qquad i = 1, \ldots, n, \tag{22}$$

where,

$$P(X) \equiv \sum_{i=1}^{q} p_i[S_i\, R_i = h_i(X)], \tag{24}$$

and where $p_i\ (S_i\ R_i)$, $i = 1, \ldots, q$ are the penalty functions associated with exceeding the environmental standards, S_i, $i = 1, \ldots, q$.

Although our optimization scheme requires only that we remove those constraints (environmental relationships) which are not of the linear form $R = AX$, we note from the formulation of the new problem, equations 23, 17, 18 and 22, that even the linear environmental models have apparently been removed (as constraints). This is optional and depends upon the model formulation and its size. If model size, the number of rows and columns, is of no consequence and if the entire management model is contained within a single linear program (LP), it is more efficient to keep the linear environmental relationships as part of the constraint set.

If, on the other hand, model size is a problem and it is desirable, as discussed above, to divide the management model up into a number of smaller LP's, disposition of the linear environmental models is not as straightforward. No matter how the larger LP is subdivided, the environmental relationships, which involve all the liquid and gaseous residuals discharges throughout the region, invariably link the smaller LP's. In this case, if the linear environmental models are retained as part of the constraint set, one of the available decomposition techniques must be employed.

---

of names. For example, Zangwill refers to this technique as "penalty" and "barrier" methods depending upon whether an optimum is approached from outside or within the feasible region. See W. I. Zangwill, *Nonlinear Programming: A Unified Approach* (Englewood Cliffs, N.J.: Prentice-Hall, Inc., 1969), Chapter 12. Fiacco and McCormick, on the other hand, refer to this as "exterior" and "interior" point methods, respectively. See A. V. Fiacco and G. P. McCormick, *Nonlinear Programming: Sequential Unconstrained Minimization Techniques* (New York: John Wiley and Sons, Inc., 1968).

The Penalty Function. The scheme which we are using to eliminate environmental relationships from the constraint set and still meet the environmental quality standards, S_i, $i = 1, \ldots, q$, is known as a penalty or exterior point method (as opposed to a barrier or interior point method). The name derives from the fact that throughout the optimization procedure we allow the vector of standards, S, equation 20, to be violated, but only at some "penalty" to the value of the objective function. The objective of the approach is to make this penalty severe enough such that at the optimum the standards will be satisfied, within some tolerance.

Because in our optimization scheme we evaluate the gradient, ∇F, at each step in the procedure, we require that the objective function, equation 23, be continuous and have continuous first derivatives. A quadratic penalty function of the following form satisfies these requirements.[27]

$$p_i(X) = \max [(h_i(X) - S_i)\ 0]^2 \qquad i = 1, \ldots, q. \tag{25}$$

For computer applications, equation 25 may be written more conveniently as,

$$p_i(X) = \left\{\frac{[h(X) - S] + |[h(X) - S]|}{2}\right\}^2 \qquad i = 1, \ldots, q, \tag{26}$$

a form which gives $p_i(X) = 0$ when $h_i(X) < S_i$.

The major difficulty with the penalty function expressed above as equation 25 is that, in general, it is not steep enough in the vicinity of the boundary (that is, the standard) and, consequently, the "unconstrained" optimum is apt to lie substantially outside the original feasible region. However, a slight modification to the P function remedies this situation. If $r > 0$, and $p_i(X) \neq 0$, the new penalty function,

$$\frac{1}{r_i} p_i(X), \tag{27}$$

approaches infinity as $r_i \rightarrow 0$. Specifying a sequence of decreasing values for r has the effect of moving the unconstrained optimum closer and closer to the boundary of the feasible region. From a computational standpoint, it is sufficient that r only be made small enough to ensure

27. Note that the second derivative of this function is also defined and that it is positive for $h(X) > S$.

that the unconstrained optimum is within a preselected distance of the boundary.[28]

This situation adds substantially to the computational requirements of our iterative optimization scheme. Not only do we have to find an optimum for the management problem given a set of penalty functions and associated parameter values, but now we have to find a new optimum for a sequence of values of r. Obviously, the fewer r's we need to use during the ascent procedure, the better off we will be. The reason for the small values of r, as noted before, is to ensure that the optimum is sufficiently close to the boundary of the feasible region. The reason for a relatively large value of r in the beginning of the ascent procedure is strictly a computational one. It is related to the efficiency of the optimization scheme employed. Rapid changes in the response surface are difficult, in general, to deal with except when the optimum is being approached and the step size is relatively short. It is difficult to know a priori what a good starting value of r would be.[29]

From an operational point of view, selection of an appropriate set of penalty functions and sequence of values for the penalty function parameter r is a real concern to us. The efficiency of the optimization scheme is directly dependent on how this is handled. We hope that with some experience with the operational behavior of a specific model, it will be possible to specify a range of values for r from which a reasonably small subset could be selected. We will investigate this question using the Delaware Valley Base Model.

The Optimization Procedure. A formal presentation of the nonlinear programming algorithm we are using to optimize a nonlinear objective function subject to a set of linear constraints has been presented elsewhere.[30] Only the essence of the scheme is repeated here. Relevant equations and expressions used for this procedure are restated, and the objective function, equation 23 is modified accordingly.[31]

28. Because at the optimum the standards are met only within some tolerance, it should be noted that $h_i X > S_i$ for some i and, hence neither the penalty, $P(X)$, nor the vector of the marginal penalties, $\partial P(X)/\partial X$, reduces to zero.

29. In addition, it should be pointed out that a relatively large value of r in the beginning of the procedure ensures that neither the value, nor the slope, of the penalty function exceeds the largest value that the computer can deal with.

30. C. S. Russell and W. O. Spofford, Jr., "A Quantitative Framework," pp. 126–37.

31. Before we proceed, it should be pointed out that other nonlinear programming algorithms do exist. Considerable progress has been made in the last decade in the development of general, nonlinear algorithms that can handle nonlinear objective functions and nonlinear constraints of both the equality and inequality type. See, in particular, F. A. Fiacco and G. P. McCormick, *Nonlinear Programming*; J. B. Rosen, "The

The optimization scheme we are using is analogous to the gradient method of nonlinear programming. The technique consists of linearizing the response surface in the vicinity of a feasible point, X^k. To do this, we construct a tangent plane at this point by employing the first two terms of a Taylor's series expansion (up to first partial derivatives). This linear approximation to the nonlinear response surface will, in general, be most accurate in the vicinity of the point X^k and less accurate as one moves farther away from this point. Because of this, a set of "artificial" bounds (constraints) is imposed on the system to restrict the selection of the next position along the response surface to that portion of the surface most closely approximated by the newly created linear surface. The selection of the appropriate set of artificial bounds is analogous to choosing a step size in other gradient methods of nonlinear programming.

Because the newly created subproblem is in a linear form, we are able to make use of standard linear programming techniques for finding a new optimal point, X^{k+1}. This point locates the maximum value of the linearized objective function within the artificially confined area of the response surface. Because, in general, the linearized surface will not match the original nonlinear surface, the original nonlinear objective function must be evaluated at this point to determine whether or not this new point, X^{k+1}, is, in fact, a better position than the previously determined one, X^k. That is, the following condition must be satisfied:

$$F(X^{k+1}) > F(X^k). \tag{28}$$

If this condition is satisfied, a new tangent plane is constructed at the point X^{k+1} and a new set of artificial bounds is placed around this point. As before, a linear programming code is employed to find a new position, X^{k+2}, which maximizes the linearized objective function, and so on until a local optimum is reached. This procedure, like all gradient methods, finds only the local optimum. If the response surface contains more than one optimum, the problem becomes one of finding the global optimum. One way of approaching this is to start the procedure at different points within the feasible region, R, where the starting points may be chosen at random.[32]

Gradient Projection Method for Nonlinear Programming, Part I: Linear Constraints." *S.I.A.M. Journal on Applied Mathematics* 8, no. 1 (1960): 181–217; and J. B. Rosen. "The Gradient Projection Method of Nonlinear Programming, Part II: Nonlinear Constraints," *S.I.A.M. Journal on Applied Mathematics* 9, no. 4 (1961): 514–32.

32. For techniques on random starts within a feasible region defined by a linear constraint set, see, P. P. Rogers, "Random Methods for Non-Convex Programming" (Ph.D. diss., Harvard University, 1966).

As we have just seen, this optimization procedure requires that we linearize the objective function at a point X^k. We do this according to the following formulation:

$$F(X^{k+1}) = \nabla F(X^k) \cdot X^{k+1} + \gamma, \tag{29}$$

where γ is a constant. Expressing our revised objective function, equation 23, along with the modification suggested by expression 27, in terms of equation 29 results in,

$$\begin{aligned} F &= \nabla f(X) \cdot X - \nabla P(X) \cdot X + \gamma \\ &= \frac{\partial f(X)}{\partial X} \cdot X - \frac{\partial P(X)}{\partial X} \cdot X + \gamma. \end{aligned} \tag{30}$$

In our residuals management problem, $\partial f(X)/\partial X$ is a vector of linear cost coefficients associated with traditional production inputs, and residuals handling, modification, and disposal activities; and $\partial P(X)/\partial X$ is a vector of marginal penalties associated with the discharge of each residual.

Given that $R = h(X)$, equation 19, we see from equation 25 that,

$$\frac{\partial P(X)}{\partial X_j} = \sum_{i=1}^{q} \frac{2}{r_i} \{\max [(R_i - S_i), 0]\} \frac{\partial R_i}{\partial X_j} \qquad j = 1, \ldots, n. \tag{31}$$

The term

$$\left(\frac{2}{r_i} \{\max [(R_i - S_i), 0]\}\right),$$

represents the slope, dp_i/dR_i, of the ith penalty function evaluated at the point R_i. The term $\partial R_i/\partial x_j$ represents the marginal response of the ith descriptor of the natural world (or ecosystem) to changes in the discharge of the jth residual. Equation 31 may be expressed more generally as,

$$\frac{\partial P(X)}{\partial X_j} = \sum_{i=1}^{q} \frac{dp_i}{dR_i} \cdot \frac{\partial R_i}{\partial X_j} \qquad j = 1, \ldots, n, \tag{32}$$

or in matrix notation as,

$$\frac{\partial P(X)}{\partial X} = \left(\frac{\partial R}{\partial X}\right)^T \cdot \frac{dp}{dR}. \tag{33}$$

For linear environmental systems, $\partial R_i/\partial X_j$ is an element of the matrix of transfer coefficients, A, when the environmental functions are expressed, linearly, as,

$$R = h(X) = A \cdot X. \tag{34}$$

Hence, for the case of linear environmental systems, the marginal penalties (equation 33) may be expressed in matrix notation as,

$$\frac{\partial P}{\partial X} = A^T \cdot \frac{dp}{dR}, \tag{35}$$

where $\partial P/\partial X$ is a vector of marginal penalties, A is a matrix of environmental transfer coefficients, and dp/dR is a vector of slopes of the penalty functions evaluated at R.

For the case of nonlinear environmental models, the situation is similar except that evaluation of the environmental response matrix, $\partial R/\partial X$, is somewhat more involved and in addition, because the response is nonlinear, it must be recomputed for each state of the natural world.

The Linearized Subproblem. Now that we have presented the essence of the optimization scheme that we are using, including a discussion of the LP subproblem which is necessary for us to both construct and solve at each step along the ascent procedure, we can restate our management problem in these terms.

$$\max \left\{ F[X^{k+1}] = \left[\nabla F[X^k] - \nabla P[X^k] \right] \cdot X^{k+1} + \gamma \right\}; \tag{36}$$

$$\text{s.t.} \quad g_i(X) = 0 \qquad i = 1, \ldots, m, \tag{17}$$

$$g_i(X) \geq 0 \qquad i = m+1, \ldots, p, \tag{18}$$

$$x_i \geq 0 \qquad i = 1, \ldots, n, \tag{22}$$

$$x_j \leq \beta_j \qquad j = 1, \ldots, s, \tag{37}$$

$$x_j \geq \alpha_j \qquad j = 1, \ldots, s, \tag{38}$$

where β_j and α_j are, respectively, upper and lower bounds on the s discharge variables at the $(k+1)$th iteration. The efficiency of the optimization scheme depends directly on how these bounds are selected. Investigation of various procedures for selecting bounds is, perhaps, the most important use of the Delaware Valley Base Model. The techniques we are

currently using will be presented later with a discussion of some results of the Base Model.

The Delaware Valley Base Model: An Illustration

In this section, we address ourselves to the computational problems discussed in the second section. To explore various solution methods and programming techniques, we constructed what we call the Delaware Valley Base Model. It is based on the Delaware Valley region in terms of geographic characteristics, economic activities, and residuals dischargers. It employs the same atmospheric dispersion and aquatic ecosystem models that the future full-scale model of this region will use. The major differences between this model and the full-scale residuals management model are the number of residuals generation and discharge activities provided with residuals management options, and the areal extent of the region considered.

The primary objective of this modeling effort is to build a computer model with all the features (hardware and software) of the proposed large scale regional model; one which can be expanded easily, but which remains small enough to experiment with programming techniques and ideas. Specifically, our aims include:

1. demonstrating the feasibility of solving a number of individual linear programs in sequence prior to entering the FORTRAN coded environmental models;
2. gaining experience with the use of complex, nonlinear ecosystem models as an integral part of the residuals management framework;
3. experimenting with the penalty function concept for meeting standards (or constraints) on ambient concentrations;
4. experimenting with various step size selectors on a reasonable size problem (136 discharge variables);
5. providing reality in terms of the Delaware Valley region.

The Delaware Valley region

The eleven county Delaware Valley region we ultimately intend to model is shown in figure 3.[33] This region consists of Bucks, Montgomery, Ches-

33. Much has been written about this particular region, especially in the water resources area. For a general discussion of water quality modeling efforts in the Delaware Estuary, see A. V. Kneese and B. T. Bower, *Managing Water Quality: Economics, Technology, Institutions* (Baltimore: Johns Hopkins Press, 1968).

ter, Delaware, and Philadelphia counties in Pennsylvania; Mercer, Burlington, Camden, Gloucester, and Salem counties in New Jersey; and New Castle county in Delaware. The rivers of interest in this region are the Schuylkill, which enters the Delaware at Philadelphia; the Delaware Estuary which runs approximately 85 miles from the head of Delaware Bay to the head of tide at Trenton, New Jersey; and a short reach of the Delaware River above Trenton. For modeling purposes, the estuary is divided into 30 sections, 16 of which are shown in figure 4.[34] For purposes of air quality management, a 10 kilometer grid is superimposed on the eleven county region. This is shown in figure 3.

The Delaware Valley region, with a 1970 population of approximately 5.5 million people,[35] is one of the most industrialized areas in the United States. For example, this region contains 7 major oil refineries, 7 steel plants, 13 paper (or pulp and paper) mills, 14 important thermal power generating facilities, numerous chemical and petrochemical plants, and 6 large municipal sewage treatment plants.[36]

The portion of the eleven county Philadelphia region which we include in the Base Model is outlined in figure 3 and shown in detail in figure 4. As depicted in these figures, the area of interest runs from Wilmington to Philadelphia. It is a rectangular area, 45 by 55 kilometers, upon which a 9 by 11 grid with equal spacings of 5 kilometers is superimposed. The river bounded by this area includes sixteen sections of the Delaware Estuary Model—sections 6 through 21.

34. These are the same 30 sections which were originally established by the Delaware Estuary Comprehensive Study (DECS) and the ones which are still being used by the Delaware River Basin Commission (DRBC) in their modeling efforts of the estuary. See Federal Water Pollution Control Administration, Delaware Estuary Comprehensive Study, and Delaware River Basin Commission, "Final Progress Report: Delaware Estuary and Bay Water Quality Sampling and Mathematical Modeling Project," May 1970.

35. This figure represents the 1970 population for the eleven county area described above; Delaware PC(VI-9), New Jersey PC(VI-32), Pennsylvania PC(VI-40). U.S. Department of Commerce, *1970 Census of Population* (Washington, D.C.: Government Printing Office, 1970).

36. Information about major dischargers in the region can be obtained from three major sources: (1) the Delaware River Basin Commission (DRBC), Trenton, N.J., for information on liquid residuals discharged to the estuary (see, for example, Delaware River Basin Commission, "Final Progress Report"; (2) Metropolitan Philadelphia Interstate Air Quality Control Region, "Inventory of Emission Sources," for information on gaseous residuals discharged within the Delaware Valley region; and (3) Greater Philadelphia Chamber of Commerce, *Business Firms Directory of Greater Philadelphia* (Philadelphia: Greater Philadelphia Chamber of Commerce, 15th edition, 1971), for addresses and employment data on plants.

Figure 3

Delaware Valley Region

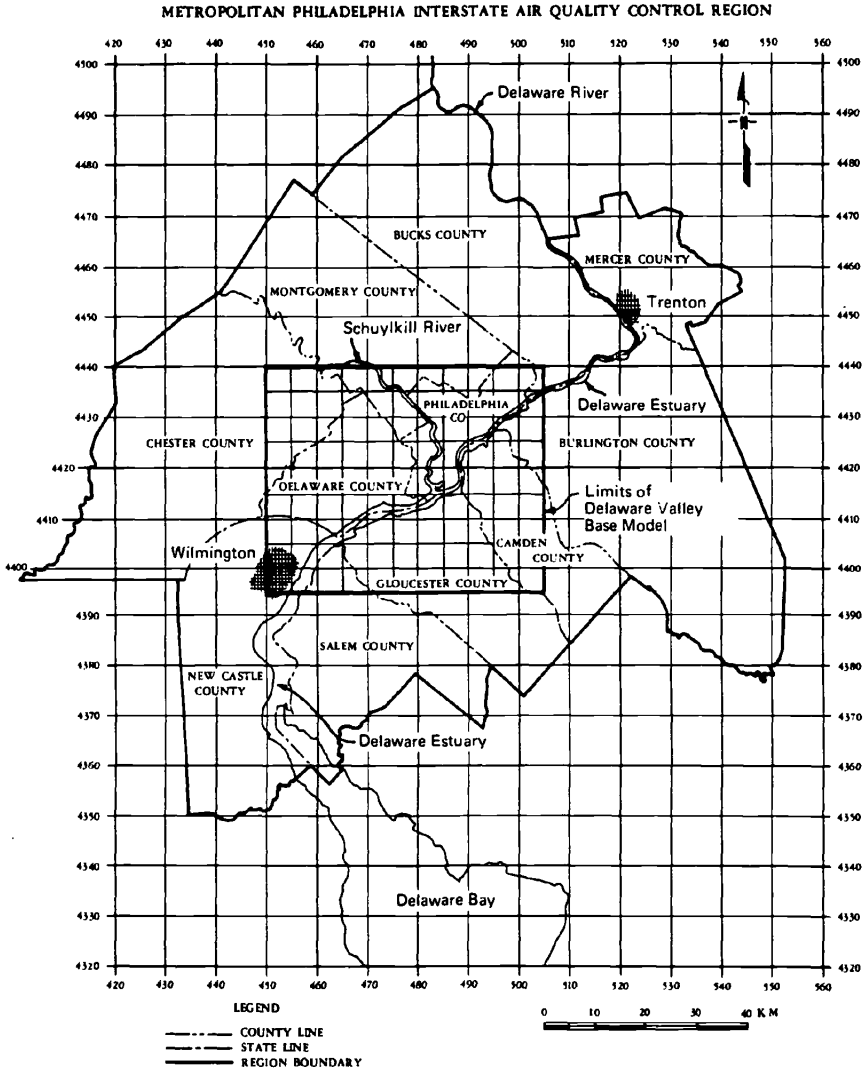

Note: The grid is in kilometers and is based on the Universal Transverse Mercator Grid System (UTM).

The area represented by the 99 (5 kilometer) grid squares has been divided into two subareas—a 50 grid square urban-suburban area, and a 49 grid square rural area. The 50 grid square urban-suburban area is employed to deal explicitly with consumer (postconsumer) residuals in the model. This area is shown shaded in figure 4.

Residuals generation and discharge activities

Ten regional activities have been modeled and provided with residuals management options; two sugar refineries, two petroleum refineries, two thermal power plants, two municipal sewage treatment plants, and two municipal incinerators. The output of the sugar refineries is refined sugar; products of the petroleum refineries include gasoline, distillate fuel, and residual fuel. Information regarding the capacities and locations of these ten activities is presented in table 1. Their locations are depicted in figure 4.[37]

As an alternative in reducing the quantities of residuals which ultimately must be handled and disposed of, outputs of these plants are allowed to vary. In addition, import and export possibilities are included in the model. The sugar refineries can shut down completely if their production levels are not constrained by employment considerations. The production of electricity within the region can be reduced, and imports used to fill up regional demand. Also, heating fuel for the region, which could be supplied by the two petroleum refineries, may be imported. The two municipal incinerators can be shut down completely.

The residuals which we consider in the base model are gaseous residuals—particulates and sulfur dioxide; liquid residuals—organic material measured by its biochemical oxygen demand (BOD), nitrogen, phosphorus, phenols (toxics) and heat; and solid residuals—furnace bottom ash, digested sludge from the municipal sewage treatment plant, wet scrubber slurries, and municipal solid wastes.

The industrial plants are assumed to be in existence so that their major features (such as the thermal efficiency of the power plants) are assumed fixed over the time span of interest, and only certain modifications may be carried out. The municipal wastewater treatment plants are assumed to have installed primary treatment only. In table 2, we summarize the various residuals management options available in the model for each type of residuals generation and discharge activity and the primary

37. It is intended that these activities be representative of industries and residuals modification facilities at these locations, but they are not necessarily accurate.

Figure 4

Map of Delaware Valley Base Model Region

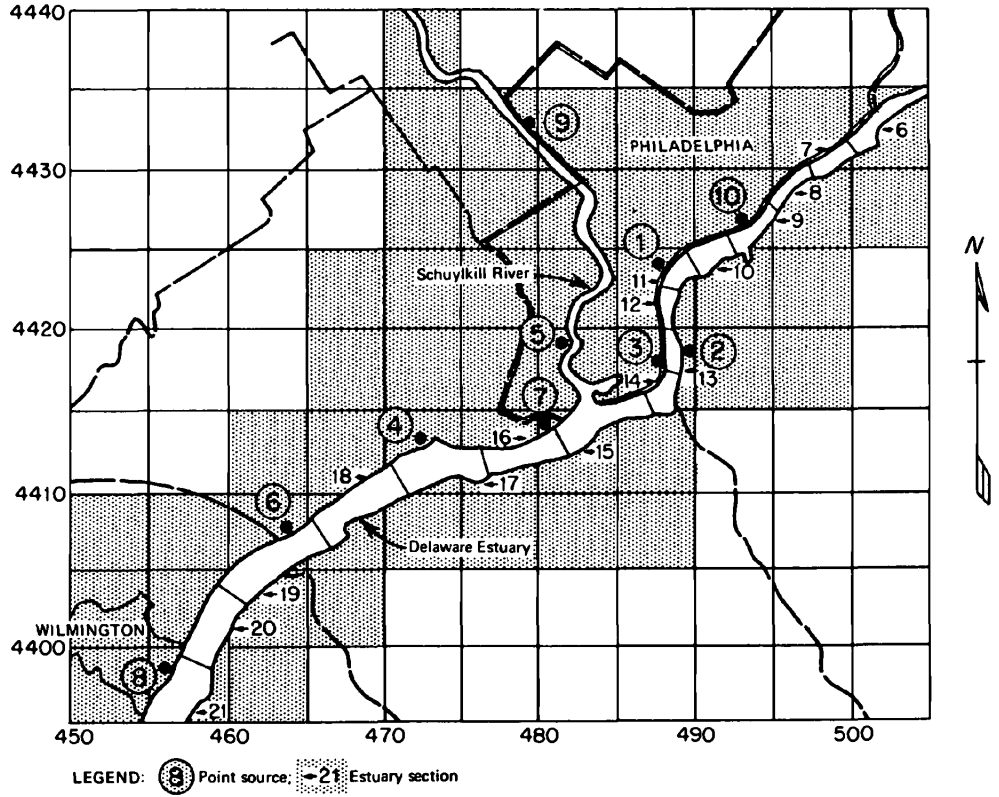

Note: The grid squares are numbered consecutively from bottom to top starting with the lower left corner and ending with the upper right corner.

TABLE 1

Delaware Valley Base Model:
Regional Residuals Generation and Discharge Activities

Source Number	Type Activity	Capacity	Location				
			City	River		UTM Coordinates	
				Reach	Mile	x	y
1	Sugar refinery	400 tons of refined sugar per day	Philadelphia, Pa.	11	101.8	488.6	4423.7
2	Sugar refinery	400 tons of refined sugar per day	Camden, N.J.	13	97.85	489.7	4418.7
3	Thermal power generating facility	345 MW	Philadelphia, Pa.	13	97.5±	488.1	4418.0
4	Thermal power generating facility	707 MW	Eddystone Borough, Pa.	17	85±	472.4	4412.2
5	Petroleum refinery	162,000 barrels of crude oil per day	Philadelphia, Pa.	15[a]	92.47[a]	482.7	4419.1
6	Petroleum refinery	172,000 barrels of crude oil per day	Marcus Hook, Pa.	19	78.95	463.8	4406.9
7	Municipal sewage treatment plant	160 mgd.	Philadelphia, SW Philadelphia, Pa.	16	90.7	481±	4413±
8	Municipal sewage treatment plant	60 mgd.	Wilmington, Delaware	21	71.3	456±	4398±
9	Municipal incinerator	565 tons per day	Philadelphia, NW Philadelphia, Pa.			479.5	4431.9
10	Municipal incinerator	298 tons per day	Philadelphia, NE Philadelphia, Pa.			493.2	4425.9

[a] This oil refinery is 2.5 miles up the Schuylkill River. The Schuylkill River enters Delaware Estuary at river mile 92.47.

TABLE 2

Residuals Management Options Available to the Various Types of Dischargers in the Region

Management Option Available	*Primary Residuals Reduced*	*Secondary Residual Generated*
Sugar Refineries		
Partial or full reuse of flume water	BOD	Sludge
Secondary and tertiary wastewater treatment	BOD	Sludge
Cooling tower(s)	Heat	Heat is rejected to the atmosphere along with water vapor
Burn lower sulfur coal	SO_2	None
Electrostatic precipitators (3 alternative efficiencies; 90, 95, 98 per cent)	Particulates (fly ash)	Bottom ash
Sludge digestion and landfill	Sludge	The secondary residual here is digested sludge at a different location
Sludge dewatering and incineration	Sludge	Particulates
Dry cyclone; 90 per cent efficiency	Particulates	Bottom ash
Petroleum Refineries		
Secondary and tertiary treatment, and various reuse alternatives (cooling tower water makeup, desalinate water, boiler feedwater)	Nitrogen Phenols BOD	Sludge
Cooling tower(s)	Heat	Heat is rejected to the atmosphere along with water vapor
Burn lower sulfur fuel	SO_2	None
Refine lower sulfur crude	SO_2	None
Sell, rather than burn, certain high sulfur products (e.g., refinery coke)	SO_2	None
Cyclone collectors on catcracker catalyst regenerator (2 efficiencies; 70, 85 per cent)	Particulates	Bottom ash
Electrostatic precipitator; 95 per cent efficiency	Particulates	Bottom ash

(continued)

TABLE 2 (*Concluded*)

Management Option Available	*Primary Residuals Reduced*	*Secondary Residual Generated*
Sell, rather than burn, high sulfur refinery coke	SO_2 Particulates	None
Sludge digestion and landfill	Sludge	The secondary residual here is digested sludge at a different location
Thermal Power Generating Plants		
Cooling tower(s)	Heat	Heat is rejected to the atmosphere along with water vapor
Burn lower sulfur coal	SO_2	None
Limestone injection-wet scrubber; 90 per cent efficiency	SO_2	Slurry
Electrostatic precipitators; 90, 95, 98 per cent efficiency	Particulates	Bottom ash
Settling pond; 90 per cent efficiency	Slurry	Solid ash
Municipal Sewage Treatment Plants		
Secondary or tertiary wastewater treatment	BOD Nitrogen Phosphorus	Sludge
Sludge digestion, drying and landfill	Sludge	The secondary residual here is digested sludge at a different location
Sludge dewatering and incineration	Sludge	Particulates, bottom ash
Dry cyclone; 80, 95 per cent efficiency	Particulates	Bottom ash
Municipal Incinerators		
Electrostatic precipitators (2 alternative efficiencies; 80, 95 per cent efficiency)	Particulates	Bottom ash

residuals which are reduced and the secondary residuals generated as a result of each of the management alternatives.[38]

38. The sugar refinery alternatives are based on information in G. O. G. Löf and A. V. Kneese, *The Economics of Water Utilization in the Beet Sugar Industry* (Washington, D.C.: Resources for the Future, 1968). The petroleum refineries are condensed versions of a model developed at RFF and described in C. S. Russell, "Residuals

As shown in table 2, there are some trade-offs among forms of residuals generated and subsequently discharged to the environment. Slurries, resulting from the removal of sulfur dioxide using limestone injection and wet scrubbing, can cause water quality problems. Particulates, resulting from the incineration of sludge, can cause air quality problems, and so on. Options do exist in the model, however, for converting the major portion of both gaseous and liquid residuals to solid residuals. Sludges are either digested, dried, and landfilled, or dewatered and incinerated, with the residue going to landfill. Furnace bottom ash is trucked directly to landfill. The slurries produced by the limestone injection-wet scrubber at the two thermal power generating facilities are lagooned on-site.

Consumption residuals

The three types of postconsumer (or consumption) residuals considered in the model are municipal sewage; particulates and sulfur dioxide from household space heating activities; and municipal solid residuals. With regard to management options in the model, we assume these residuals are generated only within the 50 grid square urban-suburban area. The quantities of sewage input to the two municipal sewage treatment plants are based on average daily BOD loadings at the Wilmington, Delaware and Philadelphia, SW wastewater treatment facilities during 1968.

Household heating requirements for the 50 grid square urban-suburban area were estimated on the basis of the Environmental Protection Agency's (EPA) inventory of gaseous emissions for this region. Sulfur dioxide emissions for each grid square, assuming the use of high sulfur distillate fuel for all household heating needs, were matched with the area source emissions.[39] This enabled us to compute the daily quantity of distillate

---

Management." The information for the electric power plants and their associated alternatives came largely from J. K. Delson and R. Frankel, "Residuals Management in the Coal-Energy Industry," unpublished manuscript; and Paul H. Cootner and G. O. G. Löf, *Water Demand for Steam Electric Generation* (Washington, D.C.: Resources for the Future, 1965). The municipal treatment plant vectors were constructed using the data compiled by Robert Smith and reported in "Costs of Conventional and Advanced Treatment of Waste Water," *Journal of the Water Pollution Control Federation* (September 1968), pp. 1546–74. The municipal incinerator characteristics are based on information contained in Combustion Engineering, Inc., *Technical Economic Study of Solid Waste Disposal Needs and Practices* Vol. IV (Rockville, Md.: United States Department of Health, Education and Welfare, Bureau of Solid Waste Management, 1969).

39. EPA's inventory of emission sources for this region consists of two kinds of sources, point and area. A point source consists of a single stack. Large sources, in terms of emissions, are usually dealt with as point sources. An area source consists of the aggregation of many smaller sources over a grid square of designated size—usually 2.5, 5.0, 10.0 or 20.0 kilometers on a side.

fuel oil required for each grid, and then assuming a heat content of 5.84 $\times 10^6$ BTU per barrel for distillate fuel oil, we were able to compute the home heating requirements (in BTU's) for each grid.

These heating requirements can be met in the model using any one of three sulfur grades of fuel oil. Sulfur dioxide emissions are assumed to be two times the sulfur weight in the fuel burned. Particulate emissions are assumed to be 0.504 pounds per barrel burned, independent of sulfur content.[40] The model can then select a policy for the sulfur content of domestic fuel oil in the urban-suburban area.

Inhabitants of the urban-suburban area have three options for dealing with their solid residuals: (1) incineration, with the residue (bottom ash) being disposed of at one of the available landfill sites; (2) disposal directly to a local landfill site with low and high compaction alternatives; and (3) disposal at a more distant, rural site with low and high compaction alternatives. Two of Philadelphia's municipal incinerators are considered in the model: Northeast, with a maximum process rate of 298 tons per day; and Northwest, with a maximum process rate of 565 tons per day.[41] We provide inputs to these incinerators by assuming an area served by each, a population for each area, and a per capita municipal solid residuals collection rate of 5.72 pounds per day.[42]

Background residuals

In order to keep the Base Model down to what we consider a reasonable size for our purposes here, only a small portion of all the residuals generation and discharge activities within the region were modeled and provided with residuals management options. Unfortunately, the resulting ambient concentrations due only to the ten point and 50 area (household space heating) sources were unreasonably low. To make the problem more interesting in terms of reality for the region, all of the remaining sources of residuals—liquid and gaseous, point and area—were then added to the model. In terms of the management framework, the newly added sources are treated as "background" inputs to the environmental models.

40. Emissions for fuel oil combustion are based on data in Environmental Protection Agency, *Compilation of Air Pollution Emission Factors* (Washington, D.C.: Environmental Protection Agency, Office of Air Programs, February 1972), pp. 1-4 to 1-7.

41. These capacities are based on EPA's inventory of emission sources for this area.

42. This is the average quantity of solid residuals collected per person in the urban United States. R. J. Black et al., *The National Solid Waste Survey: An Interim Report* (Rockville, Maryland: Department of Health, Education and Welfare, 1968), table 2.

Background inputs of residuals to the Delaware Estuary—both direct discharges and tributary loads—were estimated from data provided by the Delaware River Basin Commission (DRBC) and from surface water records published by the United States Geological Survey (USGS).[43] From the residuals discharge data on individual industrial and municipal dischargers collected by the DRBC, total inputs of nitrogen, phosphorus, and BOD were estimated for each source not already included as one of the ten activities in the Base Model. These data were averaged over the last three or four years, depending on the availability of data and the variability exhibited by the discharge stream.

Tributary loads were estimated by extrapolating gagged stream flow to the total area of each major watershed (or where tributaries were not gagged, by multiplying per area surface water runoff of an adjacent tributary by the total drainage area of the ungagged tributary), and multiplying this flow times the concentration of nitrogen, phosphorus, and BOD as reported in EPA's STORET data bank. Assumed inputs to the Delaware Estuary of BOD, nitrogen, and phosphorus have been aggregated by reach and are presented in table 3.

Background generation of sulfur dioxide and particulates was estimated as follows: All point sources in the emission source inventory, except one very large source in Grid Square 21 (which was maintained as a "background" point source) and all those treated explicitly in the Base Model as individual dischargers with management options, were aggregated over the 5 kilometer grid in which they are located. We call these "aggregated" point sources and include 44 (out of a total of 99) in the Base Model. (The remaining 55 sources were either zero or relatively small. In the latter case, they were added to the area sources of corresponding grid squares.) We assume each aggregated point source discharges through a single stack at the center of its respective grid. Stack characteristics—stack height, diameter, exit velocity, and stack exit temperature—of each aggregated point source have been specified such that the locations (heights) above the ground of the centers of mass of the

43. United States Geological Survey, *Water Resources Data for Pennsylvania: Part 1, Surface Water Records* (Washington, D.C.: United States Department of Interior, published annually); United States Geological Survey, *Water Resources Data for Maryland and Delaware: Part 1, Surface Water Records* (Washington, D.C.: United States Department of Interior, published annually); United States Geological Survey, *Water Resources Data for New Jersey: Part 1, Surface Water Records* (Washington, D.C.: United States Department of Interior, published annually). These reports are based on the water year beginning October 1 and ending September 30.

TABLE 3

Inputs of Organic Material, Nitrogen, and Phosphorus to the Delaware Estuary (pounds per day)

Estuary Sections (DRBC)	Point Sources								
	Industrial[a]			Municipal[b]			Total, Point Sources		
	BOD_5	N	P	BOD_5	N	P	BOD_5	N	P
1	0	0	0	8,065	1,210	504	8,065	1,210	504
2	2,758	2,940	807	1,668	250	104	4,426	3,190	911
3	1,111	3,169	14	2,282	342	143	3,393	3,511	157
4	1,857	328	30	0	0	0	1,857	328	30
5	2,167	202	39	644	97	40	2,811	299	79
6	0	0	0	825	124	52	825	124	52
7, 8	0	0	0	897	135	56	897	135	56
9, 10	5,530	155	31	115,525	17,329	7,220	121,055	17,484	7,251
11, 12	9,450	32	164	0	0	0	9,450	32	164
13, 14	23,659	4,380	552	141,925	21,289	8,870	165,584	25,669	9,422
15	39,973	20,385	319	4,280	642	268	44,253	21,027	587
16	25,340	5,121	176	100,564	15,085	6,285	125,904	20,206	6,461
17	63,070	17,539	455	7,093	1,064	443	70,163	18,603	898
18	16,914	5,359	277	1,803	270	113	18,717	5,629	390
19	57,036	6,800	1,311	0	0	0	57,036	6,800	1,311
20	0	0	0	673	101	42	673	101	42
21, 22	123,807	38,367	2,670	37,279	5,592	2,330	161,086	43,959	5,000
23, 24	0	0	0	335	50	21	335	50	21
25, 26	12,000	12,400	1,087	0	0	0	12,000	12,400	1,087
27, 28	0	0	0	1,365	205	85	1,365	205	85
29	0	0	0	0	0	0	0	0	0
30	0	0	0	0	0	0	0	0	0
Total	384,672	117,177	7,932	425,223	63,783	26,576	809,895	180,960	34,508

(continued)

TABLE 3 (*Concluded*)

Estuary Sections (DRBC)	*Area Sources*								
	Tributaries[e]			*Storm Water*[d]			*Totals, All Sources*		
	BOD_5	*N*	*P*	BOD_5	*N*	*P*	BOD_5	*N*	*P*
1	1,720[e]	407	514	1,360	367	103	11,100	1,980	1,120
2	2,090	931	483	0	0	0	6,520	4,120	1,390
3	0	0	0	0	0	0	3,390	3,510	157
4	182	206	238	0	0	0	2,040	534	268
5	2,000	950	425	0	0	0	4,810	1,250	504
6	4,110	1,810	724	0	0	0	4,940	1,930	776
7, 8	211	159	150	1,810	489	137	2,920	782	343
9, 10	1,880	564	439	12,960	3,499	982	136,000	21,500	8,670
11, 12	1,500	625	179	21,260	5,740	1,612	32,200	6,400	1,950
13, 14	2,790	1,785	431	9,490	2,562	719	178,000	30,000	10,600
15	17,181	4,455	5,746	18,860	5,092	1,430	80,300	30,600	7,760
16	1,430	563	299	0	0	0	127,000	20,800	6,760
17	2,470	1,737	1,000	1,950	526	148	74,600	20,900	2,050
18	879	1,012	1,037	0	0	0	19,600	6,640	1,430
19	0	0	0	0	0	0	57,000	6,800	1,310
20	0	0	0	0	0	0	673	101	42
21, 22	11,900	3,940	1,319	8,320	2,246	631	181,000	50,100	6,950
23, 24	0	0	0	0	0	0	335	50	21
25, 26	0	0	0	0	0	0	12,000	12,400	1,090
27, 28	0	0	0	0	0	0	1,370	205	85
29	0	0	0	0	0	0	0	0	0
30	0	0	0	0	0	0	0	0	0
Total	50,343	19,144	12,984	76,010	20,523	5,762	936,248	220,627	53,254

(Footnotes for Table 3)

Note: Organic material is reported in pounds of five-day biochemical oxygen demand (BOD_5) per day; nitrogen, in pounds of nitrogen (N) per day; and phosphorus, in pounds of phosphorus (P) per day.

[a] Industrial loads were estimated from unpublished data supplied by the Delaware River Basin Commission.

[b] Based on 1968 sewage treatment plant discharges, in Delaware River Basin Commission, "Final Progress Report: Delaware Estuary and Bay Water Quality Sampling and Mathematical Modeling Project," May 1970. Nitrogen and phosphorus loads were estimated from the BOD loads according to data presented in G. A. Rohlich and P. D. Uttormark, "Wastewater Treatment and Eutrophication," *Nutrients and Eutrophication* Special Symposium Volume 1 (Ann Arbor: American Society of Limnology and Oceanography, Inc., 1972).

[c] This quantity does not include inputs from the Delaware River above Trenton.

[d] BOD loads are based on 1964 stormwater overflow in R. V. Thomann, *Water Quality Management.* Nitrogen and phosphorus loads have been estimated from BOD loads according to data given in Environmental Protection Agency, *Storm Water Management Model* (Washington, D.C.: U.S. Government Printing Office, 1971), p. 180.

[e] Tributary loads are based on the low flow season averaged over a three year period.

resulting plumes of both the aggregated point source and its corresponding individual point sources (within the same grid) are the same.[44]

The area sources from the source inventory were aggregated (or disaggregated) to the 5 kilometer grid size shown in figures 3 and 4. However, we have made one modification to these sources. From the fifty area sources within our urban-suburban area, we have subtracted the gaseous emissions resulting from household space heating. Particulate and sulfur dioxide emissions from this activity are included in the residuals generation and discharge submodel, and are thus eliminated from background residuals.

We felt it was necessary to separate aggregated point sources from the area sources in the region because of the difference in effective stack heights (sum of the physical stack height plus plume rise) of these two types of sources.

The estuary and its quality

For the period October 1912 to September 1965, the mean annual flow of the Delaware River at Trenton, New Jersey, was 11,550 cubic feet per

44. This treatment does not, however, insure the same distribution of ground level concentrations. The latter is a complex nonlinear function of many arguments including effective stack height. Emissions close to the ground contribute, proportionally, more to local ambient concentration levels than do the emissions discharged higher up. The higher emissions tend to spread out over the region more. This technique was used here as an expedient measure. Consideration of each point source as a separate entity in the air dispersion model would have been an expensive proposition, and the additional accuracy would not have been warranted for the kinds of investigations we are making with our Base Model.

second (cfs). The Delaware is similar to most rivers in the northeastern United States with respect to its seasonal variation of flow: maximum mean monthly flows are experienced in March and April (20,000 to 25,000 cfs), and minimum mean monthly flows in the late summer months of August and September (about 2,500 cfs).[45]

Although a Streeter-Phelps type dissolved oxygen computer model for the Delaware Estuary is available from the DRBC, we are using the aquatic ecosystem model described in an earlier section. The DO model of this estuary used 30 reaches between Trenton, New Jersey and Liston Point, Delaware. By combining some of the shorter reaches, we have reduced this number to 22 in order to save computer time. The Delaware Valley Base Model employs 11 of the new reaches: those which lie within the area outlined in figures 3 and 4.

The atmosphere and its quality

In order to relate ambient concentrations throughout the region with gaseous residuals discharges, we needed both an air dispersion model and data on the meteorology of the Delaware Valley region. The atmospheric dispersion model which we are using in the Base Model was presented in an earlier section. Necessary data inputs to the model are a joint probability distribution for wind speed, wind direction, and atmospheric stability based on meteorological data for the Delaware Valley region. We assumed, based on EPA data, that the mean monthly maximum atmospheric mixing depth for the season of interest is 1,000 meters, and that the mean seasonal temperature and pressure for this region are 68°F (20°C) and 1,017 millibars (30.03 inches of mercury) respectively.

For purposes of the air dispersion model, each of the ten activities presented in table 1 is treated as a single point source even though most have, in fact, several stacks—one, for example, has over 40 stacks listed in the gaseous emission inventory. Where more than one stack was listed for a particular plant, a "virtual" stack was calculated on the basis of the center of mass of the individual plumes. This method of aggregation was discussed previously in this section. Stack characteristics which are required by the IPP model for computing effective stack height include physical stack height, stack diameter, and stack exit velocity and temperature of the gas.

45. United States Geological Survey, *Compilation of Records of Surface Waters of the United States, October 1961 to September 1965: Part 1–B North Atlantic Slope Basins, New York to York River* U.S. Geological Survey Water-Supply Paper 1902 (Washington, D.C.: U.S. Government Printing Office, 1970).

The computer program

The optimization algorithm we are using consists of a standard linear program (IBM's MPSX–360) and four FORTRAN coded subprograms.[46] The algorithm is set up in such a way that two, or more, linear programming problems can be solved sequentially prior to entering the four FORTRAN coded routines (to be described below). Our experience to date indicates that this is a workable arrangement as far as the MPSX software is concerned.

The ten sources in the region for which residuals management options have been provided are divided equally between two linear programs, MPSX–1 and MPSX–2. Home heating and municipal solid residuals management options for the 50 grid squares comprising the urban-suburban area have also been divided evenly between the two linear programs.

Given prices (marginal penalties) on discharges, evaluated in the environmental and constraint evaluation submodels to be discussed below, these two linear programs are solved, subject to the appropriate constraints, and the resulting residuals discharges are passed as inputs to the environmental models. See figure 5 for overall program flow.

The first of the FORTRAN subprograms, DATA SORT, is used to call the remaining three, in turn, to aggregate liquid residuals discharges entering the same reach or section of the estuary, and to assign marginal penalties to all the liquid residuals dischargers prior to returning to MPSX control for another iteration of the linear programming models.

The next FORTRAN routine, AQUATIC ECOSYSTEM, is the nonlinear aquatic ecosystem model. This model computes the steady state values for the eleven endogenous variables in each of the eleven sections of the estuary; this will be expanded to 22 reaches with the full-scale Delaware Valley Model. In addition, given penalty functions for exceeding standards on various quality measures of the natural world, this routine evaluates total penalties associated with water quality, as well as the marginal (water quality) penalties attributable to all liquid residuals dischargers in the region.

The third FORTRAN routine, ATMOSPHERIC DISPERSION, is the linear atmospheric dispersion model. In this routine, steady state levels of suspended particulates and sulfur dioxide at all ninety-nine grid

46. IBM's linear programming code was used because of its capability for handling large problems, our previous experience with it, and the availability of the IBM 360 system for our use.

Figure 5

Schematic Diagram of Program Flow: Optimization Algorithm

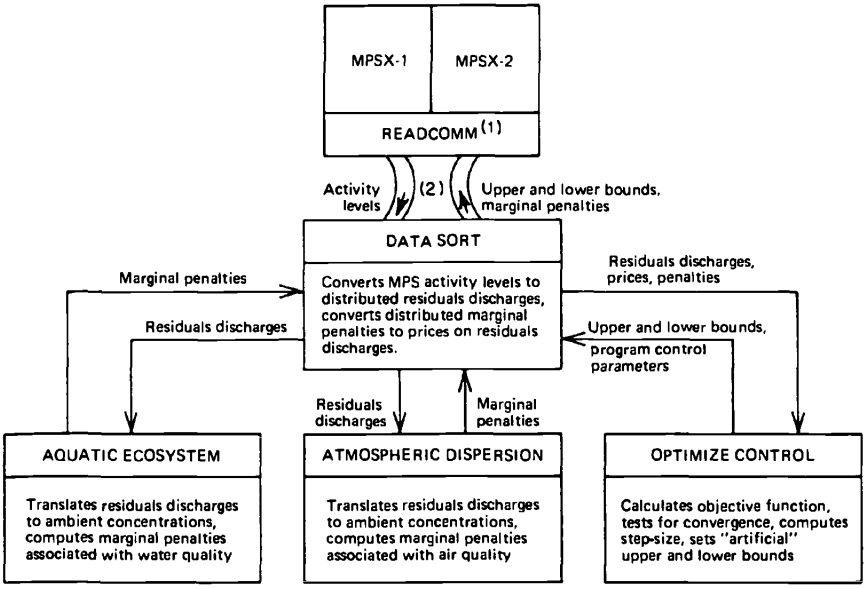

1. READCOMM is an IBM supplied FORTRAN coded subroutine which allows interaction between MPSX and other FORTRAN routines. For a description of its use see IBM Manual SH20-0960-0, *Mathematical Programming System Extended (MPSX) Read Communications Format (READCOMM) Program Description Manual* (White Plains, New York: IBM, 1971).
2. Indicates an input-output link.

points throughout the region are computed. In addition, total penalties associated with exceeding air quality standards are evaluated, and marginal (air quality) penalties are computed for each discharger in the region.

The fourth and last FORTRAN routine, OPTIMIZE CONTROL controls the allowable changes in levels of residuals discharges between successive iterations, ensures that the new solution, in terms of the objective function value, is better than the previous one, and determines when the program should terminate. Specifically, it selects the step sizes at each iteration for the 136 residuals discharge variables in the Base Model, it sets the artificial bounds (constraints) associated with each

discharge variable, and it determines when the procedure has found a local optimum.

This particular scheme of iterating between a series of linear programming problems and FORTRAN routines creates the bookkeeping problem of directing the proper set of marginal penalties and artificial upper and lower bounds to the corresponding discharge activities which are located throughout different linear programs. By properly ordering the discharge activities, and the marginal penalties and artificial bounds, the first part of the data can be read by MPSX–1 and the second part can be read by MPSX–2. This ordering takes place in the FORTRAN routine designated above as DATA SORT.

To start the iterative procedure we are currently using, we place zero prices on residuals discharge activities, and assume an initial set of values for the penalty function parameter, r. Subject to these prices and parameter values, the linear program selects "optimal" residuals discharge levels. Marginal penalties, corresponding to these discharges, are then used as prices on residuals discharges at the second iteration. This gives us a feasible starting point for the steepest ascent type scheme we employ. At the "optimum," some of the environmental standards will generally be exceeded because we are using the exterior point penalty scheme outlined above. Another set of values for the penalty function parameter, r, in equation 27 is selected such that $r^{k+1} < r^k$ for all r, and the optimization process is continued, and so on until no environmental standard, S_i, is exceeded by more than a previously specified amount. Because the response surface may be multipeaked, this procedure yields, at best, a local optimum.

One of the step size selectors which we are currently experimenting with is described in our earlier work.[47] Another, which is based on the former, will be presented in the next section where we present some results of the model. Both step size selectors require a set of ranges for the discharge variables. The feasible range for all 136 discharge variables has been estimated from a knowledge of the processes and activities employed.

For any given set of penalty function parameter, r, the program is coded to end computations when any one of the following conditions is met:

(i) the value of the objective function increases by no more than a specified amount;

(ii) the number of iterations equals the specified maximum number;

47. C. S. Russell and W. O. Spofford, Jr., "A Quantitative Framework," equation 2.1–18, p. 135.

(iii) the number of different step size sets equals the specified maximum number.

Some results

Our interest in the computational results lies primarily in the evidence they afford on the following questions:

Does the heuristic algorithm we have described appear to converge to an optimum when faced with a moderately large problem?

Is the solution obtained characterized by only small violations of the ambient environmental quality standards?

Is there evidence that the solution represents a global optimum?

What can we say about the effect of various choices for penalty-function parameters and step-size selector schemes in relation to the above questions?

Finally, is the cost of the computational exercise so high as to promise that a full-scale regional model would be simply a white elephant?

In this section, we discuss the light shed on these important questions by five separate runs of the Base Model. These runs are distinguished by the ways in which the penalty function parameters are handled, and by the step-size selector employed.

Penalty Function Parameters. Recalling that the penalty function for a particular constraint may be written $1/r \cdot p(X)$, we distinguish the following sets of values for $1/r$ in table 4:

TABLE 4

Sets of Values for 1/r

Indicator	*Penalty Function Parameter Set*			
	a	*b*	*c*	*d*
Sulfur dioxide	100	1000	100	10
Suspended particulates	10	100	10	1
Algae	10^6	10^6	10^5	10^4
Fish	10^6	10^6	10^5	10^4
Dissolved oxygen	10^6	10^6	10^5	10^4

Step-Size Selector Scheme. Three step-size selector schemes were used for the five runs of the Base Model presented in this section. For all three

step-size selectors, the initial step-size vector is the same. However, subsequent step-size vectors differ.

Given the vector of ranges, R, for all 136 discharge variables, the initial step-size vector, δ^1, is computed as follows.

$$\delta_i^1 = \frac{R_i}{10} \qquad i = 1, \ldots, 136. \tag{39}$$

For the *first* step-size selector scheme, subsequent step-size vectors are computed using the following formulation.

$$\delta^{k+1} = \frac{\delta^k}{10} \qquad k = 1, \ldots, n, \tag{40}$$

where n is the number of step-size sets.

For the *second* step-size selector scheme, step-size vectors, after the first, are computed as follows.

$$\delta^{k+1} = \frac{\delta^k}{2} \qquad k = 1, \ldots, n. \tag{41}$$

The *third* step-size selector scheme is a little more involved. Rather than shortening all elements of the step-size vector proportionately at a given step-size change (as in the case with the first two schemes), this selector reduces elements of the vector selectively. According to this scheme, at any given step-size change, an element of the step-size vector is either reduced according to equation 40, or left unchanged.

We decide whether to reduce the step-size of a particular discharge variable as follows. First, we try to assess whether the objective function decrease, which originally signaled the step-size change, is a result (on aggregate) of discharging too much or too little to the environment. We do this by comparing consecutive net benefits (the sum of the net benefits of MPSX–1 and MPSX–2 net of effluent charges), and consecutive total penalties, for the last two iterations of the management model.

If the total penalties decrease *and* at the same time the net benefits decrease, we assume that the step-size change was caused by discharging (on aggregate) too little to the environment. On the other hand, if the total penalties increase, regardless of what the net benefits do, we assume the step-size change was caused by discharging (on aggregate) too much to the environment.

Once we have established this, we must then compare, element by element, the last two sets of discharges, and the last two sets of marginal

penalties. Disposition of each element of the step-size vector is summarized as follows.

Assume discharge too much. If the discharge increases, and the marginal penalty increases, change step-size according to equation 41. For other combinations, leave step-size unchanged.

Assume discharge too little. If the discharge decreases, and the marginal penalty decreases, change step-size according to equation 41. For other combinations, leave step-size unchanged.

The Five Computer Runs. The five runs may be described as follows:

Run 1: the first step-size selector with $1/r$ set at level "a" initially and kept there for 40 iterations;[48]

Run 2: the first step-size selector with $1/r$ set at level "d" initially, kept there for 20 iterations, reset to level "a" and maintained for a further 20 iterations;

Run 3: the first step-size selector with $1/r$ set initially at level "d," changed to level "c" at the first step-size change and to level "b" with the second step-size change; the program is allowed to run out to the 40th iteration;

Run 4: the second step-size selector with $1/r$ set at level "a" initially and kept there for 40 iterations;

Run 5: the third step-size selector with $1/r$ set at level "a" initially and kept there for 40 iterations.

Ambient Standards. For each run, the ambient environmental quality standards were set as follows in table 5. The standards on the airborne residuals are much less stringent than current federal standards. They were chosen this way only because the Base Model has so few controllable sources relative to the discharges we have included as "background."[49] Standards approaching the current federal limits are infeasible in this test

48. That these runs all involve stopping after 40 iterations is simply a feature of our initially cautious experimentation. Ordinarily the algorithm itself would determine the stopping point.

49. It should also be pointed out that the air dispersion model was not calibrated prior to using it. Because of this, the computed ambient concentrations were somewhat higher than those actually observed in the Philadelphia region. If we use the calibration relationships for this model from an EPA air quality study of this region (1967–1968 data) and modify the air quality penalty function parameter sets accordingly, comparable results for our regional residuals management model could be obtained by reducing the sulfur dioxide standard in our model to $175\mu g/m^3$ ($0.416 \times 420 = 174.7$) and the suspended particulates standard to $88\mu g/m^3$ ($35 + 0.532 \times 100 = 88.2$). For these calibration relationships, see EPA, "Application of Implementation Planning Program Modeling Analysis: Metropolitan Philadelphia Interstate AQRC," Air Quality Management Branch, Applied Technology Division, Office of Air Programs, EPA, Durham, N.C., February 1972, (mimeo), figures 2 and 3.

TABLE 5

Ambient Environmental Quality Standards

Indicator	*Location*	*Standard*
Sulfur dioxide	center of each grid square	$\leq 420\ \mu g/m^3$ [b]
Suspended particulates	center of each grid square	$\leq 100\ \mu g/m^3$ [b]
Algae[a]	each reach	≤ 2.0 mg/L
Fish[a]	each reach	≥ 0.29 mg/L
Dissolved oxygen	each reach	≥ 3.0 mg/L

[a] Based on biomass concentrations in terms of the total weight of carbon, nitrogen, and phosphorus.

[b] Standards used with uncalibrated air dispersal model.

situation even under the uninteresting policy choice of shutting down all six industrial emitters and the two municipal incinerators.

Results of the Five Runs. The results of the five computer runs are presented in tables 6, 7, and 8, and in figures 6 through 10. In table 6, we show, for the five computer runs, production levels at the 40th iteration for all six industrial plants and for the two municipal incinerators. In addition, we show production levels for the first iteration (zero prices on residuals discharges) and the second iteration (maximum prices on residuals discharges). The results of the first two iterations are the same for all five runs. Also depicted in table 6 are the net benefits (sum of the individual MPSX objective function values corrected or plus for the effluent charges "paid"), the total penalties, and the objection function values at the 40th iteration for the five computer runs.

In table 7, we present for the five computer runs the discharge levels of air-borne, liquid, and solid residuals at the 40th iteration for the ten activities in the region for which residuals management options have been provided. Also shown are the residuals discharge levels for the first and second iterations (which are the same for all five runs).

Ambient concentrations which exceed the standards on the 40th iteration are displayed in table 8 for all 5 runs. There we show the standard which is violated, the amount of the violation, the location of the violation (estuary reach or air quality grid square), and the per cent violation based on the level of the standard.

Figures 6 through 10 depict a plot of the objective function value vs. iteration for the 5 runs. Also shown in these figures are the iterations

TABLE 6

Production Levels and Objective Function Values for the Five Computer Runs

Source Number	*Production Outputs and/or Inputs*	*Units*	*Iteration*		*Results at 40th Iteration for Runs*				
			1	*2*	*1*	*2*	*3*	*4*	*5*
1	Refined sugar	tons/day	400	0	58	0	0	50	0
2	Refined sugar	tons/day	400	0	107	0	41	85	5
3	Electricity	MW hrs./day	6,040	0	6,040	6,040	6,040	6,040	6,040
4	Electricity	MW hrs./day	11,550	11,550	11,550	11,550	11,550	11,550	11,550
5	Crude oil processed	bbls./day	162,000	0	129,300	147,200	145,300	129,900	144,300
	Gasoline	bbls./day	116,200	0	57,200	70,200	69,200	57,600	67,400
6	Crude oil processed	bbls./day	172,000	130,800	172,000	172,000	172,000	172,000	172,000
	Gasoline	bbls./day	123,400	53,200	99,500	123,400	123,400	105,800	121,000
9	Solid residuals processed	tons/day	495	0	506	512	516	507	505
10	Solid residuals processed	tons/day	298	0	298	298	298	298	298
Net benefits (MPSX–1 + MPSX–2)*		$ per day	492,300	69,700	297,600	335,700	317,300	306,100	329,700
Total penalties		$ per day	2,277,500	0	1,600	22,600	48,400	800	600
Objective function value		$ per day	(1,785,200)	69,700	296,000	313,100	268,900	305,300	329,100

* Effluent charges not included.

TABLE 7

Residuals Discharge Levels for the Five Computer Runs

Source Number	Residual	Units	Iteration		Results at 40th Iteration for Runs				
			1	2	1	2	3	4	5
1	Sulfur dioxide	tons/day	6.9	0	0.8	0	0	0.9	0
	Particulates	tons/day	22.1	0	1	0	0	1	0
	BOD	lbs./day	33,400	0	134	0	0	117	0
	Nitrogen	lbs./day	2,790	0	78	0	0	68	0
	Heat	10^6 BTU/day	2,380	0	0	0	0	0	0
	Total solids	tons/day	6.4	0	25.4	0	0	22.4	0
2	Sulfur dioxide	tons/day	6.9	0	1.2	0	0.7	1	0.1
	Particulates	tons/day	22	0	2.3	0	0.7	2.2	0.1
	BOD	lbs./day	33,400	0	327	0	95	198	12
	Nitrogen	lbs./day	2,790	0	180	0	56	116	7
	Heat	10^6 BTU/day	2,380	0	0	0	0	0	0
	Total solids	tons/day	6.4	0	47.2	0	18.2	37.1	2.3
3	Sulfur dioxide	tons/day	54.4	0	12.4	17.6	16.2	14.3	14.3
	Particulates	tons/day	111.7	0	23.2	22.5	21.3	23	23
	Heat	10^6 BTU/day	32,400	0	5,400	21,400	19,500	1,700	1,900
	Solids	tons/day	50.7	0	66.3	64.4	64.9	65.6	65.6
	Slurry	tons/day	0	0	173.4	152.2	157.8	165.8	165.8
4	Sulfur dioxide	tons/day	104	7.6	58.6	104	81.7	66.1	63.5
	Particulates	tons/day	309.8	9.7	140.5	275.9	226.7	168.4	184.6
	Heat	10^6 BTU/day	61,900	0	15,700	42,100	37,100	14,900	11,000
	Solids	tons/day	97	70.2	113.8	127.4	108.2	111.2	112
	Slurry	tons/day	0	285.2	187.6	0	92.2	156.6	167.2

5	Sulfur dioxide	tons/day	148.2	0	33.9	35.9	33.4	34.2	34.4
	Particulates	tons/day	13.6	0	2.4	3.5	3.5	2.4	3.2
	BOD	lbs./day	12,500	0	4,500	6,500	6,300	4,500	6,200
	Nitrogen	lbs./day	568	0	290	369	365	291	361
	Phenols	lbs./day	2,023	0	626	963	943	632	916
	Heat	10^6 BTU/day	53,200	0	740	28,400	28,000	500	0
6	Sulfur dioxide	tons/day	157	10.5	88.2	157	157	101	96
	Particulates	tons/day	14.4	1.3	7.7	14.4	14.4	9.3	12.4
	BOD	lbs./day	13,300	4,700	9,800	13,300	13,300	10,600	12,600
	Nitrogen	lbs./day	601	327	451	601	601	460	571
	Phenols	lbs./day	2,150	670	1,540	2,150	2,150	1,680	2,030
	Heat	10^6 BTU/day	56,500	0	27,200	52,700	44,900	31,700	52,900
7	Particulates	tons/day	49.4	2.1	25.4	21.3	23.2	28.6	27.3
	BOD	lbs./day	272,500	19,500	33,900	27,200	31,200	34,000	32,300
	Nitrogen	lbs./day	32,400	11,400	14,200	13,400	15,600	14,200	13,900
	Solids	tons/day	82.4	108.8	93.4	134.4	148.9	97.7	107.6
8	Particulates	tons/day	19.1	1.1	10.8	21.8	19.2	12.4	19.1
	BOD	lbs./day	105,000	7,500	27,500	72,200	58,200	31,800	105,000
	Nitrogen	lbs./day	12,500	4,400	8,300	10,800	10,200	8,900	12,500
	Solids	tons/day	31.8	55.9	58.4	36.4	42.1	54.9	31.8
9	Particulates	tons/day	9.9	0	4.9	7.4	6.6	5.7	9.5
	Solids	tons/day	123.9	0	131.9	130.8	132.8	131.3	126.8
10	Particulates	tons/day	5.9	0	2.9	4.8	4.5	3.4	5.7
	Solids	tons/day	74.5	0	77.5	75.6	76.0	77.1	74.8

TABLE 8

Violation of Ambient Standards on the 40th Iteration for the Five Computer Runs

Run	*Number of Violations*	*Indicator*	*Units*	*Location (Grid/Reach)*	*Violation, Actual*	*Violation, Per Cent of Standard*
1	4	Sulfur dioxide	$\mu g/m^3$	59	1.4	0.33
		Fish	mg/L	15	0.01	3.45
		Dissolved oxygen	mg/L	17	0.034	1.13
		Dissolved oxygen	mg/L	15	0.009	0.3
2	5	Sulfur dioxide	$\mu g/m^3$	59	5.1	1.21
		Sulfur dioxide	$\mu g/m^3$	20	0.5	0.12
		Fish	mg/L	15	0.012	4.14
		Dissolved oxygen	mg/L	15	0.13	4.33
		Dissolved oxygen	mg/L	17	0.053	1.77
3	5	Sulfur dioxide	$\mu g/m^3$	59	0.5	0.12
		Sulfur dioxide	$\mu g/m^3$	20	0.1	0.02
		Suspended particulates	$\mu g/m^3$	59	0.1	0.1
		Dissolved oxygen	mg/L	17	0.176	5.87
		Dissolved oxygen	mg/L	15	0.131	4.37
4	4	Sulfur dioxide	$\mu g/m^3$	59	1.9	0.45
		Suspended particulates	$\mu g/m^3$	58	0.2	0.2
		Fish	mg/L	15	0.001	0.34
		Dissolved oxygen	mg/L	17	0.022	0.73
5	6	Sulfur dioxide	$\mu g/m^3$	59	1.3	0.31
		Suspended particulates	$\mu g/m^3$	59	0.9	0.9
		Suspended particulates	$\mu g/m^3$	58	0.3	0.3
		Fish	mg/L	15	0.001	0.34
		Dissolved oxygen	mg/L	15	0.014	0.47
		Dissolved oxygen	mg/L	17	0.014	0.47

Figure 6

Objective Function Value vs. Iteration, Run 1

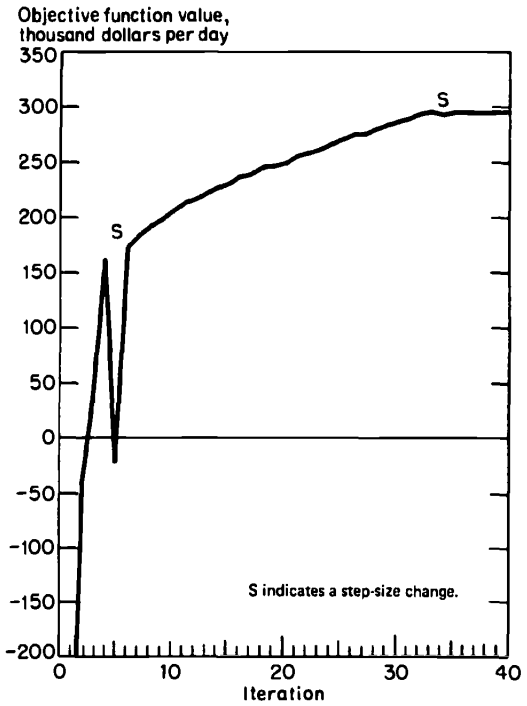

where step-size changes occurred (indicated by an "S"), as well as those where penalty function parameter sets were changed (indicated by a "P").

Discussion of the Results. The iterative optimization scheme which we are using for the Delaware Valley Base Model can find, or at least come very close to finding, a local optimum. It may take a long time, but it will eventually get there. The question we are exploring here is one of computational efficiency—which step-size scheme in combination with which sets of penalty function parameters permits climbing to the optimum the fastest. We attempt to answer this question with 5 runs of the Base Model. Caution should be exercised in comparing the results of the 5 computer runs, however, because each of the runs meets a different mix of levels of environmental quality (see table 8).

The most significant conclusions to be derived from the results of the runs follow.

Figure 7

Objective Function Value vs. Iteration, Run 2

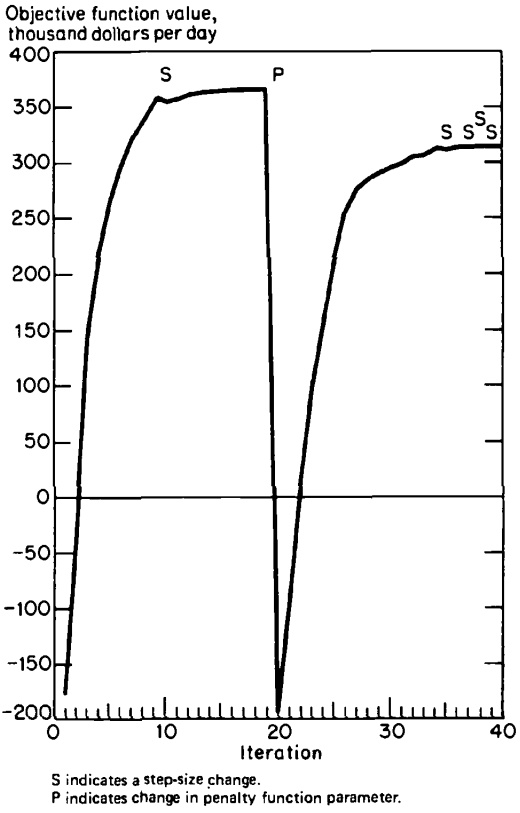

1. Technically, none of the runs have reached an optimum by the 40th iteration, although evidence suggests that Runs 2, 4, and 5 all may be quite close. Plots of typical residuals discharges throughout the iterative optimization procedure (not shown here) indicate that Runs 1 and 3 are quite far from the optimum. We might also point out that of the 5 runs, Runs 1 and 3 obtained, by the 40th iteration, the lowest objective function values (see table 6).

2. From the 5 runs presented, it is not possible to ascertain whether they are approaching the same optimum or different optima. The question of local vs. global optimum cannot be answered, but we suspect that we are dealing with a multi-peaked response surface.

Figure 8

Objective Function Value vs. Iteration, Run 3

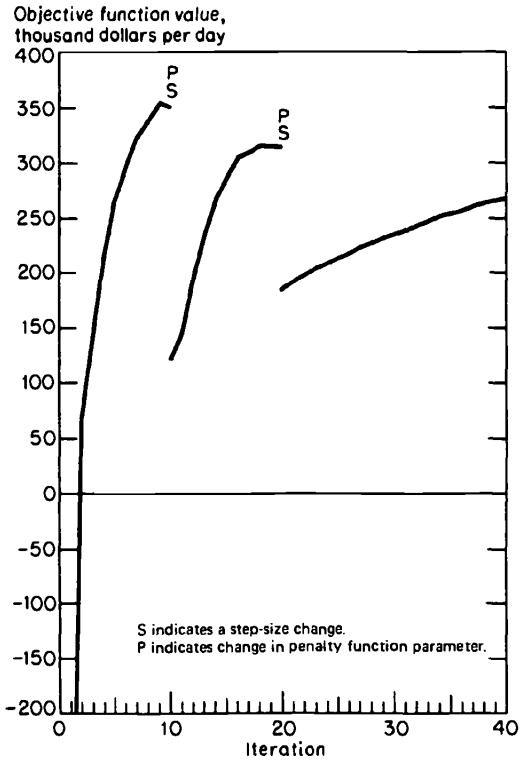

3. Of the 5 runs presented, Run 5 not only appears to be the most efficient in terms of its hill-climbing ability (see figure 10), but also it appears to have achieved the best solution by the 40th iteration. This conclusion is based on the fact that: on aggregate, Run 5 did one of the best jobs of meeting the environmental quality standards (see table 8); its objective function value is the highest of the five; its penalties are the lowest of the 5; and finally, its net benefits are the second highest of the 5 runs (see table 6).

In addition, there is evidence, albeit weak, to suggest that neither Runs 2 nor 4 are as close to the optimum as Run 5. This is shown in the liquid discharge data (BOD and nitrogen) for source 8, and the sewage treatment plant for Wilmington (see table 7). There are no standards violated

Figure 9

Objective Function Value vs. Iteration, Run 4

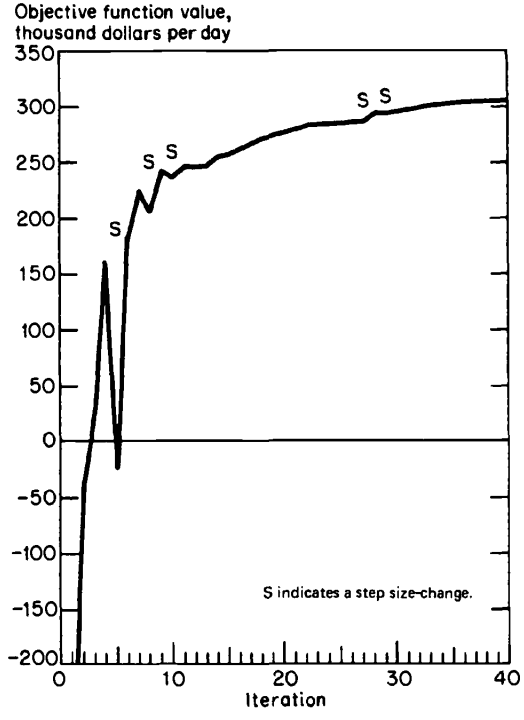

downstream of source 8, yet only in Run 5 are BOD and nitrogen discharges at their maximum levels. Both Runs 2 and 4 have quite a way to go before reaching this upper limit.

4. The penalty function scheme described in this paper was quite successful in meeting environmental quality standards (see table 8). A comparison of the sulfur dioxide violations for Run 3 (see table 8) with those of Runs 1 and 2 indicate the effect of increasing the penalty function parameter. Recall that the penalty function parameter for Run 3, on the 40th iteration, is ten times that of either Runs 1 or 2.

5. The third step-size scheme with no change in penalty function parameter, Run 5, appears to be superior to the other schemes which were tested. However, the step-size choice continues to remain the weakest part of our optimization technique and there is much room for improvement.

Figure 10

Objective Function Value vs. Iteration, Run 5

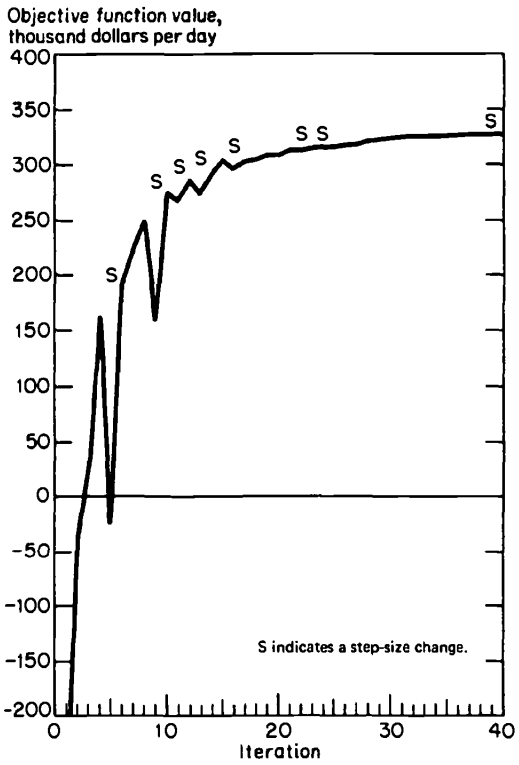

6. Examination of residuals discharge data throughout the iterative optimization scheme (not shown here) suggests that the response surface in the vicinity of the optimum is relatively flat, at least inside the feasible region, and that substantial tradeoffs among different residuals and different dischargers, in many cases, are taking place. Examination of the net benefit function values on the 40th iteration (see table 6) together with the rather significant variations in discharge data for this iteration (see table 7) indicates the same thing. That is, there are wide ranges of discharges among the 5 runs, yet the net benefits vary only a few per cent (maximum of between nine and ten per cent for Runs 2, 4, and 5). This difference is even more pronounced when net benefits and discharge data for Runs 2 and 5 are compared.

Costs of a Computer Run. What can we say about costs, and the prospects for running a much larger version of this model within a reasonable computer budget? To place this question into some perspective, we summarize, in the following table, the size, in terms of the number of rows and columns of the LP submodels, of the Delaware Valley Base Model.

TABLE 9

Size of LP Submodel
—Delaware Valley Base Model—

	MPSX–1	*MPSX–2*	*Total*
Variables (columns)			
Total	320	320	640
Residuals discharges	68	68	136
Constraints (rows)			
Total	178	178	356
Equality	135	135	270
Inequality	43	43	86
Step-size bounds (rows)	136	136	272

The aquatic ecosystem model comprises 11 reaches of the estuary, with each reach requiring a set of 11 nonlinear differential equations. This results in a total of 121 equations for describing the aquatic system within the limits of the Base Model. The number of reaches to be considered in the full-scale Delaware Valley model will be expanded to 22. This amounts to an eventual doubling of the size of the aquatic ecosystem model.

The air dispersion model includes 11 point sources, 44 "aggregated" point sources, and 99 area sources for a total of 154 sources. Ambient concentrations are computed for 99 receptor locations. In order to predict the ambient concentrations of sulfur dioxide and suspended particulates, one 154 by 99 element matrix of transfer coefficients is required for each residual. It is difficult to estimate at this time how many sources will eventually be considered in the full-scale Delaware Valley model, but the receptor locations probably will not be increased by more than 50 per cent.

The cost per iteration of the Base Model is about $8.70 [49] for the cur-

50. This is based on a CPU running time of 0.42 minutes per iteration on an IBM 360–65 with a core requirement of 300 K.

rent computer program. A large fraction of this cost is for internal input-output time (65 per cent); a smaller fraction for central processing unit (CPU) time (26 per cent); and the smallest part for computer printout (9 per cent). It goes without saying that there are large savings possible if the number of iterations required to reach, or at least come close to, an optimum can be reduced through more efficient means of selecting step-sizes and sequences of penalty function parameter sets. It is our plan to continue to search for more efficient combinations. But the cost information above also implies that there are potentially large savings from reducing the quantity of information which must be stored and transferred in and out of core with the various submodels each iteration.

One possibility for saving internal input-output time would be to include the linear air dispersion model within the constraint set of the residuals generation and discharge LP submodel. This would involve adding 198 rows (99 for ambient concentrations of sulfur dioxide, and another 99 for suspended particulates) to the Base Model, but could reduce the internal I–0 time by, perhaps, ten per cent. In addition, it would probably improve the efficiency of the iterative optimization scheme since in this case penalty functions would be needed only for standards on algae, fish, and, dissolved oxygen. However, a major drawback of this approach is that individual (unlinked) LP submodels, in general, would not be possible with this scheme and some form of decomposition would be necessary to deal with model size.

We are currently searching for other ways to reduce internal input-output time, even at the expense of increasing CPU time, and plan to investigate some of them with the Delaware Valley Base Model.

As for cost increases with model size, the problem appears serious, but certainly not serious enough to abandon the project at this point. It appears to us that, currently, most of the CPU time is used for sorting both discharges and marginal penalties, for controling the program flow, and for the solution (including penalties and marginal penalties) of the environmental submodels. CPU costs, however, should not increase in proportion to the increase in the number of discharges because many of the computations necessary to set marginal penalties at each iteration are already being done in the Base Model.

On the other hand, the input-output time required for the transfer of residuals discharges, marginal penalties, and the changing upper and lower bounds on each discharge will increase linearly with the number of discharges. This is a problem, and one which lends little hope for improvement as long as nonlinear simulation models of the natural world are included within the optimization framework. But as noted above, this

time requirement for the Delaware Valley model could be reduced substantially as long as we remain with linear atmospheric dispersion models and include them as part of the constraint set of the LP submodel.

The costs of solving the MPSX packages themselves appear to be a relatively small part of the problem, at least at this time, but it also appears that it will be more efficient to have a few large MPSX LP's rather than many smaller ones. The primary reason for this is that each MPSX module has certain minimum internal input-output requirements independent of its size.

On balance, we are hopeful that the model size can be increased considerably with a substantially less than proportional increase in the cost of a computer run.

Summary and Conclusions

In the Introduction, we presented, very briefly, our regional residuals management framework and indicated that two didactic applications to hypothetical regions had previously been made and are reported elsewhere. In the first application, both demand functions and economic damage functions were assumed, and the institutional framework envisioned was a regional management authority. In the second application, the model was expanded to provide information on the socio-geographic distribution of costs and benefits associated with meeting *different levels* of environmental quality. The institutional framework envisioned for selecting levels of environmental quality was a legislative body. We further indicated in the Introduction that the main purpose of this paper is to explore the computational problems associated with scaling up from small didactic models to a large scale regional application.

We have discussed, in some detail, the problem of model size and ways to cope with it, and also presented a way to include nonlinear models of the natural world within an optimization framework. As a means for testing our ideas and techniques for dealing with some of the computational problems presented we constructed a relatively simple residuals management model of the Delaware Valley region. This application is referred to by us as the Delaware Valley Base Model. Results of five computer runs indicate that our optimization scheme, which allows for the inclusion of nonlinear models of the natural world, is operationally feasible, but that more work is necessary to improve the computational efficiency of the procedure so that computer costs can be reduced.

In conclusion, there is no question that the management model we are experimenting with is more complex than previous residuals management models, and that it requires a substantial amount of computer resources. Three questions we face and hope to answer as a result of our research efforts on large scale regional residuals management models are as follows.

1. Is it necessary to include all forms of residuals within a single computational framework? In principle, it is necessary, but in practice, does the additional simultaneous information on air and water quality, and on the generation of solid residuals, warrant the added effort and expense?

2. Nontreatment management options such as changes in production processes, raw material input mixes, and product specifications, and by-product production and recycling, have all been shown (under some circumstances) to be less costly alternatives to reducing, or modifying, residuals than waste treatment alternatives, but can we consider even a portion of these alternatives for a large complex region consisting of many different types of economic activities without (a) exceeding the computational reliability of present day computer hardware and software; and (b) exceeding the computer budget of a waste management agency given that the first condition could even be satisfied?

3. Models of aquatic ecosystems are able to provide additional useful information for making public policy on environmental resources, but (a) can we incorporate these models within an optimizing framework without completely expending the computer budget of a regional residuals management agency; and (b) is it necessary to go to nonlinear ecosystem models given all the other uncertainties, in both data and model components, in the other parts of the regional residuals management model?

Our research to date indicates that the inclusion of nonlinear models of the natural world within an optimization framework is expensive, but not unreasonably so. Whether it is necessary to include these nonlinear models, or whether it would suffice to employ linear models of the natural world, we cannot say at this time. We are currently in the process of exploring this question.

Regarding the second question, given that the intent of our regional residuals management modeling effort is to be able to generate distributional information on costs, benefits, and environmental quality for a wide range of alternative management strategies for meeting ambient environmental quality standards, a priori elimination of management options, in many cases, would be a difficult, and at best arbitrary, task. Our

research thus far has shown that nontreatment alternatives are frequently less costly than the more traditional abatement alternatives, but even more important to us, it has shown that, in most cases, a priori selection of alternatives for least-cost solutions is not possible because of all the links—both market and nonmarket—which exist for any complex situation. In order to be able to provide as many management options in the regional model as seems desirable for the Delaware Valley region, we plan to continue to search for ways of coping with model size.

Finally, as far as the first question is concerned, none of us will know the answer to this until someone or some group tries an integrated approach to residuals management modeling and compares its output with that of other kinds of residuals management models.

COMMENT

J. Hayden Boyd, The Ohio State University

The paper by Spofford, Russell and Kelly is a progress report on the Delaware Base Model, an ongoing, large scale modeling effort at Resources for the Future. An earlier paper by Russell and Spofford also reported on the progress of the model, and gives much useful background on its conception and planning.[1] Two sorts of questions naturally arise in examining a paper such as this. First, what are the technical goals of the modeling effort which it describes, and how well have these technical goals been achieved? Second, what are the nontechnical goals of the model? That is, how do the authors plan to use the model to increase the value of our environment, as we consume its waste disposal and environmental quality services?

The authors wish to advance the state of the art in several ways. First, several waterborne, airborne, and solid wastes, and several aspects of water and air quality, are considered simultaneously. For waterborne wastes and water quality, the model contains 11 endogenous variables, including 5 residuals (N, Ph, BOD, phenols, heat), and 3 quality parameters (algae, fish, DO). For air, sulfur dioxide and particulates are modeled. Solid wastes such as furnace bottom ash, digested sludge from municipal sewage plants, wet scrubber slurries, and municipal solid wastes are also included. Second, the environment is modeled using a nonlinear

1. Clifford S. Russell and Walter O. Spofford, Jr., "A Quantitative Framework for Residuals Management Decisions," *Environmental Quality Analysis,* Allen V. Kneese and Blair T. Bower (Baltimore: The Johns Hopkins Press, 1972).

aquatic model and a linear atmospheric model. The linear atmospheric model is simple and straightforward, while the aquatic model is quite complex and expensive to compute. The authors claim that "nonlinear representations of the natural world increase the complexity and number of calculations ncessary for each iteration, but also they increase both the realism and predictive capability of the model." They present no evidence to support this conclusion.

Third, and significantly, the authors are modeling production modifications in ten waste dischargers. They are allowing the amount of waste generated to be modified as a pollution control strategy, and they are also looking at trade-offs among wastes. For example, sewage treatment produces sludge which when burned produces particulates which can be precipitated and carried to a landfill.

This is an impressive and ambitious list of technical features, and represents an advance in the state of the art. The talents of competent systems analysts and the capacities of a large digital computer are challenged by an effort of this scale. It is important, however, to keep perspective by considering what the authors were not able to do. They did not consider alternatives such as regional treatment plants or bypass piping to move discharge points elsewhere in the estuary. These kinds of alternatives typically reduce the cost of achieving a given level of water quality. They did not consider expansion and contraction of the scale of current waste generators, not to mention the options of entry and exit from the region. Stochastic elements or cyclical variations in waste generation or in environmental reaction were not considered. Environmental modifications (for example, low flow augmentation) are not studied in the present paper, although the earlier paper by Russell and Spofford indicated that at one time there were plans to include environmental modifications. Explicit demands for environmental quality (the so-called "damage functions"), also included in the earlier plans, have apparently been dropped.

How well have the limited, albeit ambitious technical goals been achieved? As the authors are the first to acknowledge, the results are not yet in. The model is complex enough that analytical solution is impossible, and gradient methods must be used. There is little theory about which hill climber is fastest or how to speed up convergence. I have no doubt, however, that the authors will succeed in solving these knotty technical problems.

In its present form, the model is far short of fulfilling the authors' technical goals. Only ten waste dischargers are modeled explicitly. "Background" residuals seem to dominate environmental quality. Shutting down the ten modeled sources would not meet federal standards, but

adding more sources explicitly will only add to computational difficulties.

There were also some expositional difficulties. I had difficulty following the description of the industry submodels. Several alternative techniques for shoehorning the industry models into the overall model were described, but it was not indicated which one was actually being used. The inference is that the outputs of plants are allowed to vary as a pollution control strategy, since imports seem to be a substitute for the plants' outputs in some cases. But, I was not able to find an explicit statement about how plant outputs do, in fact, vary as part of the overall strategy of pollution abatement. It is not clear exactly how the choice between low sulfur fuel and high sulfur fuel for household heating is being handled. Finally, the explanation of why penalty payments on discharges ought not to appear in the overall objective function was hard to follow. The reason is of course that while a payment for this service is a cost to the firm, it is not a cost to society, as it represents a rent on the scarce environmental resource.

The particular gradient method which the authors have selected deserves further comment. The "penalty functions" may appear to be surrogate prices on effluents, but they are not. The marginal penalties go to zero as the environmental quality constraints are met. In fact, it is this very feature which leads to computational difficulties, because the algorithm gets progressively "lazier" in the vicinity of the feasible region. The $1/r$ parameters are designed to give an extra push as the feasible region is approached. The behavior of the algorithm reflects the implicit economic assumption made by this and other such programming models: that the demand for the constraint is zero elastic. In other words, improvements in environmental quality beyond the standard are valueless, while violations of the standard cost infinite consumer surplus, so that the standards will be met (if at all possible), whatever the cost. The penalty or exterior point method of programming temporarily relaxes the constraint as a computational expedient, but the end result is a maximum which lies in the "feasible" region.

Environmental quality constraints are always somewhat arbitrary, particularly here where background residuals are so high that the federal standards can't be met in any case. It would be better to retain in the model an explicit demand function for environmental quality, even if it has to be assumed at this stage of our knowledge. The aquatic and atmospheric models could then compute the marginal opportunity cost of the pollutant disposal services of the environment and pass them to the industry submodels as effluent charges. These effluent charges would not

go to zero as the feasible region (satisfying a set of arbitrary environmental quality constraints) is approached. An explicit quality demand function may help to speed up convergence of the model. Even if not, it would certainly make easier the economic interpretation of the model's solutions.

The overall goal of the modeling effort is to help rationalize the management of the environment. Neither in this paper, nor in the earlier paper by Russell and Spofford, is the institutional framework within which this model is to be used discussed in detail. But the authors do envision a regional management authority with power to set effluent charges or standards. In other words, the authority has ownership rights in the environment, and allocates its capacity among various outputs, including the absorption of various kinds of wastes and the provision of various aspects of environmental quality. The model is to be a staff tool for such an authority.

Regional environmental management models are not new, but as yet they have been little used to guide the actual allocation of the environment's scarce services. Any attempt by a regional authority to ration the environment's waste absorption services affects the wealth of dischargers, giving incentives to combine with others to use the political mechanism to influence the effluent charges or standards. If the regional management authority uses effluent charges, raising charges to reduce effluents results in consumer surplus losses to the dischargers. If discharge standards are used, pollution reductions cause changes in both consumer surplus and the implicit expenditure rectangle associated with the consumption of the environment's waste disposal service. If a systems analysis technique such as the Delaware Base Model is used to get a "least cost" solution, there may be severe transfers of relative wealth among dischargers, as some are cut back more than others. Next year's model, which considers a different set of alternatives, may lead to quite a different distribution of relative wealth. The political system seems to resist procedures which lead to capricious wealth transfers. In the absence of institutions for side payments, it deems to serve up uniform treatment standards (e.g., secondary treatment for all, best available technology), even when lower (total) cost solutions have already been documented.

According to the authors, the purpose of this paper is to answer the question, "Have we developed a mildly interesting academic curiosity or a potentially useful management tool?" One infers that the authors believe that if the computer program should converge fairly rapidly then the Delaware Base Model is a useful management tool. Yet, in the ab-

sence of innovative institutions to cope with the wealth transfer problem, it seems doubtful that this model will actually lower the social cost of a clean environment.

One object of this modeling effort is to "find a 'best' set of policy instruments (charges and limits) for imposition by . . . the authority." [2] The authors wish to shed light on the choice among alternative allocating institutions: standards (allocation by fiat) vs. use of the price mechanism; effluent charges vs. discharge rights (perhaps auctioned, perhaps transferable among polluters).

I wonder if a super systems analysis, modeling simultaneously as many relevant aspects of the universe as can be put into the computer, is really the vehicle to answer such questions. Disaggregated, piecemeal research would seem to be indicated, rather than a large scale computer model, to inform the choice among alternative allocating institutions. We need evidence on information requirements and administrative costs. What about the ability of alternative institutions to cope with stochastic and cyclical elements in waste generation and in the environment? What about the effects of the choice of institutions on the growth and decay of various industries? What about the ability of institutions to cope with the entry and exit problem?

The Delaware Base Model surely has social utility even if, because of the expense of its computation, it turns out to be a "mildly interesting academic curiosity." But it would seem that much information beyond that capable of being generated by such large scale models is required to aid in our search for a better environment at least social cost.

2. Russell and Spofford, "A Quantitative Framework," p. 119.