

REPORT DOCUMENTATION PAGE

AFRL-SR-BL-TR-98-

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

0750

ering
on of
Suite

1. AGENCY USE ONLY (Leave Blank)	2. REPORT DATE December, 1993	3. REPORT TYPE AND DATES COVERED Final	
4. TITLE AND SUBTITLE Optimal Conjunctive Management of Coupled Surface Water and Groundwater Systems Using Gradient Dynamic Programming		5. FUNDING NUMBERS	
6. AUTHORS Charles Russell Philbrick, Jr.			
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Stanford University		8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) AFOSR/NI 4040 Fairfax Dr, Suite 500 Arlington, VA 22203-1613		10. SPONSORING/MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES			
12a. DISTRIBUTION AVAILABILITY STATEMENT Approved for Public Release		12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) See Attachment			
14. SUBJECT TERMS		15. NUMBER OF PAGES	
		16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT UL

DTIC QUALITY INSPECTED 8

OPTIMAL CONJUNCTIVE MANAGEMENT OF COUPLED SURFACE WATER
AND GROUNDWATER SYSTEMS USING GRADIENT DYNAMIC
PROGRAMMING

A THESIS
SUBMITTED TO THE DEPARTMENT OF GEOLOGICAL AND
ENVIRONMENTAL SCIENCES
AND THE COMMITTEE ON GRADUATE STUDIES
OF STANFORD UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
MASTER OF SCIENCE

BY
Charles Russell Philbrick, Jr.
December 1993


19981202 050

Approved for the Department:

A handwritten signature in cursive script, appearing to read "Steve Lindick", written over a horizontal line.

Advisor

Approved for the University Committee on Graduate Studies:

A handwritten signature in cursive script, appearing to read "George F. W. ...", written over a horizontal line.

ABSTRACT

To meet demands for limited water resources, managers of water supply systems are considering innovative management methods. "Conjunctive use"—the coordinated management of surface water and groundwater—is frequently espoused as a method for addressing these demands by improving the reliability and efficiency of water management activities.

Water management is often heuristic, based on prior experience managing existing systems. Because conjunctive-use is relatively new, we do not have heuristic rules for managing the allocation of water between surface water reservoirs and aquifers.

This thesis presents an analytic procedure to derive rules for managing a conjunctive-use system. In this thesis, system management refers to the specification of optimal controls: decisions regarding the release and allocation of water. The procedure is applied to a proposed conjunctive-use system of practical interest, where the facilities have pre-determined characteristics and where control decisions are affected by uncertain autocorrelated inputs. Optimal decisions are determined by applying the optimization method of gradient dynamic programming to a numerical model of the conjunctive-use system.

The resulting optimal decisions demonstrate that application of the analytic procedure can greatly improve the management of a conjunctive-use system. These decisions are significantly more efficient than those that result from application of heuristic rules. Also, the analytic procedure incorporates the effect of autocorrelation, and the resulting optimal decisions demonstrate that autocorrelation is important in the control of conjunctive-use systems because of the different capabilities and constraints that surface reservoir storage and aquifer storage present.

ACKNOWLEDGMENTS

Research and studies in support of this work have been funded by the United States Air Force Laboratory Graduate Fellowship Program sponsored by Air Force Office of Scientific Research (AFOSR) and conducted by the Southeastern Center for Electrical Engineering Education (SCEEE). In addition, computer resources utilized by this work have been provided by National Science Foundation Grant BSC-8957186 and the Hewlett Packard Foundation.

TABLE OF CONTENTS

	<u>Page</u>
Abstract.....	iii
Acknowledgments.....	iv
Table of Contents.....	v
 I. Introduction	
A. MOTIVATION FOR THIS WORK: PLANNED CONJUNCTIVE USE BY A WATER SUPPLY AGENCY.....	1
B. DEVELOPMENTS IN RESERVOIR MANAGEMENT AND CONJUNCTIVE USE.....	3
C. OBJECTIVES OF THIS WORK.....	5
 II. Optimal Control of Stochastic Autocorrelated Stream Flows in a Conjunctively Managed Surface Water and Groundwater System	
A. INTRODUCTION.....	7
1. Differences Between Reservoir and Aquifer Storage	
2. The Importance of Autocorrelation in Optimal Control	
3. Literature Review of Conjunctive-Use Control Problems	
B. OPTIMAL CONTROL.....	14
1. Concept	
2. General Model	
C. PROBLEM DESCRIPTION.....	19
1. System Description	
2. Optimal Control Problem	
3. Stochastic Model of Stream-Flow Time Series	
D. SOLVING THE CONJUNCTIVE-USE PROBLEM.....	34
1. Method	
2. Interpretation	
3. Simple Illustration	
4. Stochastic Versus Deterministic	
5. Convergence	

E. RESULTS.....	44
1. Efficient Allocation	
2. Effect of Autocorrelation	
3. Optimization and Simulation	
4. Possible Simplifications	
F. CONCLUSIONS.....	55
 III. Modeling Issues	
A. QUANTIFYING SYSTEM CHARACTERISTICS.....	57
1. The Objective Function	
2. The Probability Distribution of Stochastic Inputs	
3. Incorporating Uncertainty	
4. Modeling Stochastic Variables	
5. Incorporating Trends, Discount Rates, and Seasonality	
B. COMPROMISES.....	68
1. Selecting State Variables	
2. Length of a Stage	
 IV. Solution of Autocorrelated Stochastic Dynamic Control Problems by Gradient Dynamic Programming	
A. INTRODUCTION.....	71
1. Discrete Dynamic Programming	
2. Gradient Dynamic Programming	
3. Extension to Include Autocorrelation	
B. PROBLEM SPECIFICATION.....	76
1. Modeling a Control Problem	
2. Solving a Control Problem	
C. THE GDP METHOD.....	79
1. Notation	
2. Method	
3. Algorithm	
D. COMMENTS AND CONCLUSIONS.....	87
1. Efficiency	
2. Convergence	
3. Conclusions	

V. Discussion and Concluding Remarks	
A. EXTENSIONS OF THIS ANALYSIS.....	90
B. CONCLUSIONS.....	91
C. DEFINITION OF VARIABLES.....	91
Appendix A.....	94
CONSTRAINING THE OPTIMIZATION PROBLEM	
1. Formulation of the Constraints	
2. Types of Constraints	
Appendix B.....	98
FUNCTION APPROXIMATION THROUGH HERMITE INTERPOLATION	
Appendix C.....	102
DEVELOPMENT OF DERIVATIVES	
1. Derivatives of an Objective Function $f(\mathbf{x}, \mathbf{u}, \mathbf{w}, \mathbf{y})$	
2. Approximation by Discrete Realizations	
3. Solving Gradients du^*/dx	
List of References.....	109

Tables.....	
1. State variables needed	30
2. Constraints on state and decision variables	30
3. Stream-flow parameters	33
Illustrations.....	
II-1. Schematic representation of the conjunctive-use system	29
II-2. Hydrograph of historical monthly stream flow	31
II-3. Fitted monthly flow quantiles	31
II-4. Autocorrelation and partial autocorrelation of monthly stream flows	32
II-5. Recursive development of control solution	40
II-6. Optimal decisions applied to the first recursion	41
II-7. Optimal decisions applied to the second recursion	42
II-8. Solution convergence	43
II-9. Expected cost of water rationing	50
II-10. Minimum cost allocation of stored water	51
II-11. Minimum cost allocation of stored water, Δ prior stream flow	52
II-12. Supply decision	53
II-13. Groundwater allocation decision	54
III-1. Regression of stochastic variables on state variables	67
IV-4. Flow chart for the GDP algorithm	86
B-1. Illustration of a hyper cube of 3 dimensions	101

I. Introduction

Increasing demands on limited water resources are leading managers of water supply systems to search for system designs and control methods that are efficient and innovative. "Conjunctive use"—the coordinated management of surface water and groundwater—has received much attention as an affordable and environmentally sound method for enhancing water supply systems. Increasingly, conjunctive use is proposed as a solution to water supply problems (*Sax et al.*, 1991; *EBMUD*, 1992; *Western Water*, 1993), and it is likely that we soon will see implementation of large-scale conjunctive-use projects. As yet, however, there have been few projects implemented and few guidelines exist for their management.

Because of the lack of heuristic or analytic guidelines, system managers cannot anticipate how they should optimally control new or modified systems that include both surface water and groundwater. In particular, the different capabilities and constraints of surface-water-reservoir (hereafter abbreviated to "reservoir") and aquifer storage result in optimal controls that cannot generally be found by common sense or heuristic approaches. A purpose of this thesis is to illustrate how we can find optimal controls and how the different capabilities and constraints of reservoir and aquifer storage affect optimal control. I will accomplish this by modeling and solving the control decisions that are optimal for a conjunctive-use system.

A. MOTIVATION FOR THIS WORK:

PLANNED CONJUNCTIVE USE BY A WATER SUPPLY AGENCY

Managers of water supply systems often hold two extreme though conflicting views when contemplating the addition of aquifer storage to an existing reservoir system. On one hand, they often desire to view aquifer storage in terms of an equivalent reservoir. On the other hand, as *Lettenmaier and Burges* [1979] note, they tend to view aquifer storage as a back-up, to be used in times of plenty only after reservoirs are full and in times of shortage only after reservoirs are empty. Both extreme views fail to recognize the different capabilities and limitations of reservoir and aquifer storage.

In reality, optimal control of conjunctive-use systems is far from either of these extreme views. This thesis draws attention to some of the fundamental

differences between aquifer and reservoir storage, as well as the benefits of integrated management. If we are to operate a system efficiently and reliably, management rules cannot be based on notions that groundwater provides only a backup supply or that aquifer storage is equivalent to reservoir storage.

I base my comparison of the two storage mechanisms on the analysis of a hypothetical conjunctive-use system. This hypothetical system is derived from the East Bay Municipal Utilities District (EBMUD) system in California. EBMUD supplies water to many communities along the eastern shore of the central and southern San Francisco Bay. EBMUD anticipates a growing demand for water resources in its district; and it has been studying ways of improving its ability to meet these demands.

Until recently, large water supply agencies such as EBMUD relied on the expansion of reservoir storage to meet their growing needs. Additional storage allows agencies to make more effective use of fluctuating and uncertain stream flows. To make effective use of its fluctuating and uncertain water source of the Mokelumne River, EBMUD initially planned to develop the Buckhorn Reservoir. This reservoir would have added capacity to an extensive reservoir system consisting of five reservoirs in EBMUD's service district and two large reservoirs in the Sierra Nevada Mountain Range.

However, two events have thwarted EBMUD's efforts to develop Buckhorn Reservoir. The first was a set of legal suits brought against EBMUD in 1989. Five groups charged the District with failure to consider adequately other options: federal and state law requires consideration of impacts and alternatives in an Environmental Impact Statement /Report (EIS/EIR) before construction of a Reservoir. The second was the 1992 elections for EBMUD's board of directors. The elections produced a public debate on the construction of Buckhorn Reservoir and resulted in the replacement of four of the five board members with individuals who campaigned against the Reservoir.

As a result, Buckhorn Reservoir is no longer a serious option for enhancing the EBMUD system. Conjunctive use, on the other hand, has become a serious option; and the revised EIS/EIR identifies it as the "environmentally superior alternative" among the methods considered [EBMUD, 1992]. In 1993, the EBMUD board of directors endorsed conjunctive use as the primary method for enhancing its system.

EBMUD's study of conjunctive use in its revised EIS/EIR originally motivated this work. Though the District initially expressed little desire to

implement a conjunctive-use program, it has now become the primary method for enhancing the future reliability and yield of its water supply. It is likely that other utilities facing needs similar to EBMUD will implement conjunctive-use methods.

Though development of new or modified conjunctive-use systems is likely to increase, there are, at present, no established guidelines that can aid planners in even the most preliminary evaluation of benefits. I intend this work to provide a procedure for efficient control of a conjunctive-use system, and thus clarify how planners and managers may evaluate and operate such systems.

B. DEVELOPMENTS IN RESERVOIR MANAGEMENT AND CONJUNCTIVE USE

Much of our understanding and practice of water management is the result of many years experience regulating water systems. Based on this experience, managers have developed operating rules that minimize risk and expected costs arising from water shortages and floods, and that maximize the benefit of hydropower generation.

For many systems, the years of operating experience have produced rules that are highly efficient; thus analytic optimization may provide little additional efficiency. For example, optimization of an existing hydropower system by *Kelman et al.* [1990] improved efficiency only slightly (though slight improvements may amount to a considerable financial benefit). Pre-existing rules for this system already were highly efficient after years of experience gained by trial-and-error. However, *Kelman et al.* developed their rules with much less effort and without subjecting the system to sub-optimal control while gaining operating experience. Moreover, they were able to develop a variety of other rules for different constraints—an ability that is valuable for system planning when constraints are flexible.

Though many existing systems are efficiently managed, in recent years two developments have increased the need for research in quantitative methods for water systems management. The first development is that some of the objectives of reservoir management have changed. Many system managers must now include the "costs" and "benefits"—though ill-defined—of qualities such as impacts on fisheries, riparian ecosystems, scenery, recreation, and water quality. Assuming that these values are quantifiable, system managers face the problem

of incorporating these qualities in new operating rules. The second development is that reservoir construction has become very difficult. There are fewer remaining sites for reservoir construction and the public is less accepting of this as a solution for addressing society's increasing competition for limited water supplies. As *Lettenmaier and Burges* [1979] point out, "The best available sites for surface storage facilities have been used and environmental considerations have eliminated some remaining potential sites." As a result, new and innovative management techniques are being developed. One of the most popular is the integration of surface and groundwater resources in a single "conjunctive-use" management system.

The concept of conjunctive use has been around for many years. The 1957 *California Water Plan* pointed out that in the Central Valley underground storage within 200 feet of land surface is about six times larger than feasible surface storage. Also, research applied to the conjunctive-use problem began years ago: one of the first quantitative descriptions of a conjunctive-use management problem was by *Buras* in 1963. Although the potential benefits of using the large capacities of aquifer storage have long been recognized, few projects have been undertaken and existing literature on the management of conjunctive-use systems is limited.

Management of conjunctive-use systems presents an optimal control problem that can be divided into two phases: modeling/simulation and optimization. While the modeling and simulation of water supply systems is common, optimization is used infrequently. Unfortunately, available optimization methods have been limited in their ability to accurately address realistic problems, and no single method is generally applicable. Because of this limitation, most problems thus far have been simplified excessively. *Burges and Maknoon* [1975] note that, "A feature of nearly all the literature is the assumption that one or several parameters or variables dominate the problem at hand." Consequently, most literature provides little practical guidance to system managers seeking to plan and control real systems. Not only have past efforts ignored important physical characteristics of the conjunctive-use problem, details about important non-physical characteristics arising from economic and legal factors have been absent or lacking, especially when compared with the physical and hydrologic detail normally included [*Burges and Maknoon*, 1975].

C. OBJECTIVES OF THIS WORK:

This thesis will focus on the effect that differences in storage capacity and transfer rates have on optimal control of reservoir and aquifer storage. The storage capacity of reservoirs is generally more limited than that of aquifers, but the transfer rate of water in and out can be very large. In contrast, the storage capacity of aquifers is generally much larger, but the transfer rate is limited.

As a result, efficient conjunctive-management practices are different from management practices for either surface or subsurface water supplies alone. Reservoir storage has greater ability to respond to needs during short-term but severe droughts, while aquifer storage has greater capacity to meet water needs during long-term but moderate droughts. Efficient management of a system containing both reservoir and aquifer storage requires the optimal use of the different capabilities and limitations of each.

This work presents two distinct developments: (1) optimal control of a conjunctive-use system that does not rely on heuristic groundwater allocation rules and that incorporates the influence of uncertain autocorrelated inputs, and (2) an optimization method that is capable of solving this control problem while allowing sufficient complexity to provide a solution of practical value. Optimal control—release and allocation decisions that best achieve the objectives of a system—enables us to observe general characteristics of control decisions aimed at allocating water between reservoir and aquifer storage and at hedging against uncertainty.

To solve a conjunctive-use problem, I apply a modified gradient dynamic programming (GDP) method. GDP is a dynamic programming method developed by *Foufoula-Georgiou and Kitanidis* [1988] to solve dynamic control problems of greater complexity than previously possible. In this thesis, I extend the GDP method to dynamic control problems that contain autocorrelated stochastic inputs. An autocorrelated stochastic input, such as a time-series of stream flows, has an expected future value that depends on its prior values. Our knowledge of these prior values can significantly affect our control decisions. Taking advantage of GDP's ability to solve more complex problems—and its extended ability to include autocorrelation—we can solve conjunctive-use control problems with fewer simplifying assumptions.

The solution of the conjunctive-use control problem presented in this work illustrates two important observations about the optimal control of such systems. The first observation is that appropriate control of aquifer storage should not rely on heuristically or arbitrarily defined operating rules. This is a weakness of prior research on optimal control of conjunctive-use systems: prior research generally has assigned aquifer storage a back-up role that fails to recognize the different capabilities and limitations of reservoir and aquifer storage. By constraining management decisions, this back-up role has resulted in control solutions that are less efficient than those that could have been obtained by the approach presented here. The second observation is that autocorrelation can significantly affect the optimal control of conjunctive-use systems because of the different characteristics of reservoir and aquifer storage.

Prior studies of conjunctive use systems are deficient in several respects: (1) they lack models that are realistic enough to solve practical management problems, and (2) they lack an optimization approach that is capable of solving any but the most simplified problems. The two developments of this thesis partially address these two deficiencies. Chapter II considers characteristics of optimally—vice heuristically—derived system control and the effect of autocorrelation on optimal control. Chapter IV presents an extended GDP method that can allow us to solve optimal control problems that are more complex than those solvable by other methods. Chapters III considers additional aspects of system modeling, and Chapter V presents conclusions, practical extensions of this work, and a glossary of terms.

II. Optimal Control of Stochastic Autocorrelated Stream Flows in a Conjunctively Managed Surface Water and Groundwater System

A. INTRODUCTION

This chapter illustrates and solves a conjunctive-use control problem. The conjunctive-use model is based on an existing water supply system that currently uses reservoirs for system storage, but which soon may include aquifer storage facilities. The optimal control solution is developed without the use of heuristic control rules and incorporates the effects of uncertain autocorrelated inputs.

Through this example, we can observe significant differences in the characteristics of reservoir and aquifer storage and the effect that these differences have on optimal control. This clarifies the benefits of conjunctively managing surface and subsurface water resources, thus allowing system planners to better anticipate the costs and benefits of proposed conjunctive-use systems. Optimal management of conjunctive-use systems takes advantage of the strengths of both reservoir and aquifer storage mechanisms, achieving greater system efficiency and reliability than alternate systems using either mechanism in isolation.

In this section, I discuss the unique character of the conjunctive-use problem by focusing on the different capabilities and constraints of reservoir and aquifer storage. In addition, I discuss the influence that autocorrelation can have on optimal control, an influence that can amplify the differences between reservoir and aquifer storage. Such influences have generally been ignored by prior conjunctive-use research as reflected in the literature review in the last part of this section.

The following sections of this chapter develop and solve a conjunctive-use problem. Section B formulates the generic control problem that Section C applies to a specific system model. Section D discusses the solution method. Sections E and F present results and conclusions drawn from the solution of the specific problem.

1. Differences Between Reservoir and Aquifer Storage

In a dynamic system, current decisions influence future decisions by altering future constraints on those decisions and their expected cost. As a result, system operators make control decisions that balance current performance against future performance. When operators sacrifice current performance for

improvements in the future, they make decisions that "hedge". For example, a rational operator will accept rationing of current water supplies if the cost is more than offset by savings from less severe rationing in the future.

Uncertainty plays a crucial role in producing control decisions in dynamic system. When control decisions are regulated by uncertain inputs (such as from stream flow or imperfect measurements of the state of a system), it is not possible to identify, in advance, a single control schedule that optimally balances current and future performance through time. Instead, system operators must determine control decisions in "real time"—they must continuously update control decisions based on realized and predicted values of inputs. Because realized and predicted values of inputs improve estimates of future system performance, operators can improve control decisions that balance current performance against expected future performance.

Stream flow and many other hydrologic phenomena are, by their nature, uncertain. Conjunctive-use systems will generally depend on a stream-flow source; and, as a result, the uncertainty inherent in this source will result in optimal control decisions that hedge supply and allocation decisions. However, unlike control systems consisting only of surface water reservoirs, the mix of storage mechanisms in conjunctive-use systems results in less obvious hedging policies. The different constraints on the use of surface and subsurface storage suggest that the optimal control of the different components may be quite different.

Two fundamental differences between subsurface and reservoir storage are their potential size and rate of response. Available subsurface storage frequently exceeds available surface storage by several orders of magnitude. On the other hand, filling and draining of reservoirs can occur at rapid rates, while recharge and pumping of groundwater can be significantly constrained. These fundamental differences indicate that optimal control of reservoirs and aquifers should be quite different: while we might store water in reservoirs for use during the next season or year, we might store water in the subsurface for a drought several years in the future. *Lettenmaier and Burges [1979]* compare the performance of groundwater and surface water systems by stating, "In contrast to the rather long-term failure modes encountered in excessive reliance on groundwater supplies, shorter scale (e.g., annual or seasonal) failures may result from exclusive use of surface supplies. The difference in time scales results

because typical surface storage reservoir volumes are much smaller compared to abstractions than are groundwater supplies."

Because of these differences, water supply systems should be designed to employ the strengths of both mechanisms, resulting in greater reliability and lower cost than a system using either method in isolation. As *Burges and Maknoon* [1975] point out, "Whenever multiple sources of water with different characteristics, as is the case with groundwater and surface water systems, are available, it may be possible to develop an operating strategy which exploits the different characteristics of the sources." *Willis and Yeh* [1987] recognized that, "By controlling the total water resources of a region, conjunctive use planning can increase the efficiency, reliability, and cost-effectiveness of water use, particularly in river basins with spatial or temporal imbalances in water demands and natural supplies."

In addition, because many of the largest water supply systems have relied entirely on reservoir storage, there may be many easy opportunities to utilize aquifer storage. Initial efforts to employ aquifer storage may be far cheaper than expansion of reservoir storage. *Lettenmaier and Burges* [1979] found that, for their model, developing aquifer storage to buffer against variations in stream flow was about an order of magnitude cheaper than developing surface storage.

In spite of the potential benefits, few water supply systems have implemented conjunctive-use programs, and few guidelines exist for their operation. Because of limited experience and the complexity of these systems, efficient control of new or modified systems is difficult to determine before actual operation. However, development of heuristic operating rules may result in unnecessarily large costs early in a project's life.

2. The Importance of Autocorrelation in Optimal Control

Because of the role that uncertainty plays in producing efficient control decisions, accurate representation of this uncertainty is essential. One should incorporate any information that can improve estimates of uncertain inputs. At time scales used in real-time system control, autocorrelation of stream-flow time series is high; and, as a result, measurements of prior stream-flows are essential to accurate estimation of future flows. Other information may also contribute to estimates of future stream flow.

Incorporating autocorrelation is particularly important for our understanding of the conjunctive-use problem. Because the response

characteristics of reservoir and aquifer storage are different, optimal control of a conjunctive-use system varies with the degree of input autocorrelation. We can consider a conjunctive-use system to be analogous to an electrical circuit that consists of components whose impedances (i.e., resistances to flows of electrons) depend on the frequency of an applied signal. The current flow and stored energy in each component vary with different applied frequencies. Likewise, a conjunctive-use system consists of storage "components" whose optimal control decisions depend on the varying stream-flow "signal". The optimal release and water level of each storage mechanism vary with different stream-flow autocorrelation.

Changes in the stream-flow signal are most significant when they affect optimal decisions that ration and allocate water supplies. By incorporating autocorrelation in the model of a conjunctive-use system, we can determine efficient hedging policies that are more efficient and useful for practical problems. This is an important development of this work, made possible by my concurrent development of the optimization tool presented in Chapter IV.

3. Literature Review of Conjunctive-Use Control Problems

Nearly all early conjunctive-use literature reviewed by *Burges and Maknoon* [1975] assumes that one or several parameters dominate the problem at hand. This assumption was necessary because of optimization methods available at the time were limited. Though methods and computational abilities have advanced since, more recent problems remain excessively simplified.

Among the earliest optimal control studies of conjunctive use, *Buras* [1963] models a steady-state two-component surface/groundwater system using a coarse discrete dynamic programming (DDP) approach. He models the aquifer storage mechanism essentially the same as an equivalent reservoir. Pumping costs—though included—are linear with pumping rate. The problem does not consider the quadratic effect of draw down on cost or rate limits on pumping or recharge.

Buras [1972] continues this early work, proposing but not solving a stochastic four-dimensional sequential decision problem that includes a reservoir, aquifer storage, salinity, and prior month's stream flow. To reduce memory requirements, he recommends a polynomial return function to approximate the minimum expected cost of future system operations. A heuristic rule specifies control of groundwater and restricts it to a backup role

used only after the reservoir is empty or full. A constant polynomial function of pumping rate approximates the changing pumping costs resulting from groundwater draw down.

Gal [1979] focuses on the autocorrelation of inputs by proposing a conjunctive-use model for Lake Kinneret and two aquifers that make up a major part of Israel's system. *Gal* incorporates two prior stream-flow measurements using a second-order auto-regressive model. Because this model poses a difficult optimal control problem with five dimensions, he simplifies the model and solves an optimal control problem with three-dimensions that does not include the aquifers. Also, *Gal's* proposed problem does not include constraints on pumping rate and assumes linear rationing and pumping costs.

Lettenmaier and Burges [1979] investigate the ranges of storage capacity, groundwater pumping and recharge parameters for which conjunctive-use appears feasible. Stochastic stream flows are incorporated by a Monte-Carlo analysis of synthetic stream-flow series. A heuristic rule is used to specify control of groundwater, restricting it to a backup role. *Lettenmaier and Burges* convert the effective storage benefit obtained by aquifer storage to an equivalent surface reservoir, and observe that effective storage decreases as pumping and recharge capacities are reduced. However, they note that the effect is not substantial until these capacities are reduced below about one-half the annual system demand.

Bogle and O'Sullivan [1979] study the effect of hedging in a system consisting of a reservoir, uncertain stream flow, and an alternative source. The alternative source provides water at a fixed price per unit volume which could approximate a groundwater source under certain conditions. Though the alternative source can provide water with certainty, it alone cannot meet the whole demand, so shortfall is possible. The authors specify the supply rule for the alternate source explicitly as the difference between demand and release, up to the alternative's maximum capacity. Costs arise from use of the alternate source and penalty cost for low storage and shortfall. Hedging occurs in decisions that attempt to maintain storage levels, though the approach is heuristic.

Existing conjunctive-use solutions that consider stochastic inputs generally rely on heuristic control policies. Heuristic control policies used by *Buras* [1972], *Lettenmaier and Burges* [1979], and *Bogle and O'Sullivan* [1979] fail to

consider the possible benefits of more flexible operating guidelines that might anticipate and mitigate future rationing costs.

Willis and Yeh [1987] presented two complex conjunctive-use models that incorporate many realistic problem characteristics. Their first model is a two-component surface/groundwater system that uses a response matrix approach to incorporate groundwater heads, allowing the use of quadratic pumping costs and bounds on heads. Their second model is a three-period linear problem containing 49 constraints, 24 decision variables, and 30 state variables using the simplex method and based on the Arkansas River Valley in Colorado. Both problems, however, avoid stochastic inputs.

The conjunctive-use problem encompasses two areas of research in the modeling and management of hydrological systems: surface reservoir control and aquifer management. Uncertainty from stream-flow inputs presents the largest challenge in determining control decisions for reservoir control problems. In contrast, parameter uncertainty is the principle focus of many aquifer management problems. Though input uncertainty and parameter uncertainty both affect optimal control decisions, they are very different. Input uncertainty is dynamic and makes it impossible to identify a single optimal control schedule in advance. In contrast, parameter uncertainty is not dynamic: once quantified, this uncertainty does not evolve with the system. For a problem that considers only parameter uncertainty and not input uncertainty, we can still identify a single optimal control schedule.

Researchers in reservoir control and aquifer management have used two different approaches in solving problems. On one hand, reservoir control problems reduce all systems to simple lumped parameter models that do not consider smaller scale variability in the phenomena described. On the other hand, aquifer management problems generally ignore input uncertainty and the effect this uncertainty has on the control of dynamic systems.

These two approaches have developed because of limitations on available optimization techniques. By ignoring input uncertainty, we can solve problems by a variety of optimization techniques and can model complex systems using many variables [*Willis and Yeh*, 1987]. By including uncertainty, we generally are restricted to models with few variables and to the optimization technique of discrete dynamic programming (DDP: not the same as differential dynamic programming) that quickly becomes computationally infeasible with additional variables [*Yakowitz*, 1982; *Yeh*, 1985]. In general, aquifer management problems

ignore input uncertainty [Gorelick, 1983; Casola et al., 1986; Gorelick, 1987; Jones et al., 1987; Wagner and Gorelick, 1987, 1989; Willis and Yeh, 1987; Yazicigil and Rasheeduddin, 1987; Georgakakos and Vlatza, 1991; Wagner et al., 1992; Tiedeman and Gorelick, 1993] and most reservoir control problems rely on simplified models [Yakowitz, 1982; Yeh, 1985; Saad and Turgeon, 1988; Kelman et al., 1990; Johnson et al., 1991; Karamouz and Vasiliadis, 1992; Saad et al., 1992]. Maddock [1974] developed operating rules for a stream-aquifer system that incorporates uncertain demand, though the problem solves only a single optimal control schedule based on an equivalent deterministic formulation.

In this work, I use an optimization method that is more flexible than methods available for past work. This allows me to solve a conjunctive-use problem characterized by stochastic inputs with greater accuracy and less simplification than required by past studies. Chapter IV develops and discusses the optimization method of Gradient Dynamic Programming.

Though a basic conjunctive-use system is relatively simple, prior efforts have not addressed two important characteristics of these systems. The first is the effect that autocorrelation can have on optimal control. This research is the first to consider this effect without heuristic constraints on system control. This allows us to observe control decisions that hedge and take advantage of the distinct characteristics of reservoir and aquifer storage. The second is the effect that distributed characteristics of groundwater systems can have on pumping costs. Though not addressed in this thesis, I hope to address this second characteristic in later work by linking distributed parameter models of aquifer management with the discrete parameter models of reservoir control, perhaps in an approach similar to Willis and Yeh, 1987.

B. OPTIMAL CONTROL

1. Concept

There are a variety of criteria that can be used to evaluate the performance of control decisions applied to the management of a conjunctive-use system. These criteria include maximum revenue, minimum risk, or minimum deviation from a standard. In this problem, performance is measured by the cost of control decisions; and optimal control decisions are those that minimize expected cost. Cost is a convenient and common basis for valuing an objective and is especially convenient when combining different objectives in a multi-objective problem. For simplicity, I assume that all costs are known and quantified. Chapter III discusses the formulation of an objective function when costs are ambiguous or difficult to quantify.

The manager of a water system must periodically make two types of decisions that determine how water is allocated. The first is how to allocate water in "space"—i.e., between the storage mechanisms and uses—while obeying system constraints and anticipating the future stresses on the system. The second is how to allocate water in time; when and how much to allocate to demands to best achieve benefits or avoid costs.

Decisions that allocate water in space and in time are both affected by uncertainty. A manager's need to anticipate the system's response without certain knowledge of future hydrological conditions leads to the decision to hedge. For a water supply system, a manager might hedge by rationing supplies to reduce the chance of more severe shortages in the future. For a flood control problem, a manager might hedge by allowing limited short term flood damage to reduce the chance of more severe flooding later. These are examples of how a manager may hedge allocation decisions in time. Though incurring extra costs for recharge and pumping, a manager might hedge by storing water in the ground to allow greater opportunity to retain flood-water in a reservoir, simultaneously maximizing opportunity to use flood water and minimize flood risks. This is an example of how a manager might hedge allocation decisions in space.

A manager makes periodic "real-time" control decisions based on currently available information about the system. System operators apply these decisions for a discrete period, or a "stage", of operations until a manager verifies

or updates control decisions based on changes in system information. These decisions are based on various unchanging characteristics of the system that constrain control decisions (such as reservoir sizes, aqueduct capacities, statutory requirements, contracts, etc.) and on various changing characteristics that affect how one operates within the constraints (such as current reservoir levels, demands, stream flows, snow-pack accumulation, weather patterns, time of the year, etc.). Though both the changing and unchanging characteristics of the system affect the optimum control decisions, only those characteristics that change over time produce changes in the control decisions. These changing characteristics of the system describe the state of the system—given the state of the system, one knows everything that is available to help identify optimal control decisions. Both the state of the system and the control decisions can be summarized by variables that, respectively, are called state variables and decision variables. The goal of the optimal control problem is to determine the optimal value of decision variables as a function of state variables.

2. General Model

To determine optimal controls for a real system, we may describe the system by an abstract mathematical model that describes the characteristics, constraints, and objectives of system operation. In the work presented here I assume that a system can be represented by a mathematical model that lumps system characteristics into a limited number of variables.

The following generic model uses standard "state-space" notation frequently used in reservoir management modeling.

a. Variables

Control decisions are described by a vector of m "decision variables" $\mathbf{u}$. These decisions may include how much water to release to meet needs, how to allocate available supplies to available storage, and so forth. A manager makes these decisions based on current values of changing characteristics of the system defining the initial "state" of the system and described by a vector or n "state variables" $\mathbf{x}$. The system evolves to state $\mathbf{y}$ at the end of a stage with the application of control decisions and realizations of stochastic inputs. Stochastic inputs are represented by a vector of r "stochastic variables" $\mathbf{w}$. They can include uncertain stream flows, measurement errors, or other system

characteristics that prevent exact prediction of y . This evolution is modeled by a transition function

$$y = T(x, u, w) .$$

For simplicity, parameters x , u , w , and y are with reference to the current stage and thus are not indexed. A state y at the end of a stage t is the state x at the beginning of the next stage $t+1$; i.e., $y_t \equiv x_{t+1}$.

Though the actual values of w are unknown until the end of the stage, I assume that the multivariate probability distribution of w is known and composed of independent stochastic variables. When w contains variables that are autocorrelated, such as time series of stream flows, realized values of w from prior stages condition the probability distribution. These prior values also affect the expected cost of control decisions and the optimal solution $u^*(x)$; and as important information on the state of the system, they are included as state variables. With this extension of x , we can specify the distribution of w by an unchanging or "stationary" model $w(x)$ that depends only on the state vector.

b. Constraints

Without constraints, variables describing a system can take on values that are not feasible for a real system. Some constraints exist due to relationships between variables such as those defined above; some exist due to physical or practical limits. Frequently, constraints exist on the potential states of the system and on control decisions. These are bounds on state and decision variables:

$$x_{\min} \leq x \leq x_{\max} , \quad y_{\min} \leq y \leq y_{\max} , \text{ and } u_{\min} \leq u \leq u_{\max} .$$

Feasibility of initial state x is ensured by choosing only feasible values; thus constraints on x do not enter into the optimization method discussed below. Other constraints may exist to define variables needed to model a system but are represented by equality constraints in the transition equation or directly in the objective function.

Constraints must be respected by control decisions. The standard form of inequality constraints is

$$g(x, u) \leq d(x)$$

where g and d are scalar functions. Using the transition equation $y = T(x, u, w)$, we can formulate constraints that do not depend on y . I assume that decisions must be feasible regardless of any realized value of stochastic input w ; thus, with only bounds as constraints on variables, we can more simply represent them by

$$g^T u \leq d(x)$$

where g is a vector of scalar coefficients. Appendix A provides further discussion of the development of constraints for reservoir control problems.

c. Objective Function

The measure of performance of the decisions is an objective function f that measures expected cost of control decisions. Costs might include losses from hedging (e.g., rationing and minor short-term flood or environmental damage), operational costs/benefits (e.g., allocation expenses such as from pumping and hydropower generation) or a variety of qualities that produce extra-market costs or benefits (e.g., water quality, recreation, and wildlife habitat). The objective function can be dependent on control decisions, state of the system, and uncertain inputs. Given the "optimal decisions" $u^* = u^*(x)$ that minimize costs given a state x , the optimum value of the objective function is

$$f^*(x) = \min_u \{f(x, u, w, y)\} = f(x, u^*, w, y)$$

For optimal control of reservoir systems, *Foufoula-Georgiou and Kitanidis* [1988] propose a general objective function composed of two costs: those incurred in the current stage and those expected in following stages. Current cost results directly from current control decisions and may also depend on the state of the system. Expected cost also results indirectly from current control decisions because of the effect that current decisions have on limiting or changing future decisions and states of the system. The objective function is represented by

$$f = E_w \{C(x, u, w) + F(y)\}$$

where C is the current cost function and F is the expected cost function (also known as the "cost-to-go" from state y). Because w is uncertain, y , F , and C are also uncertain. The objective function is the expected cost calculated over the

distribution of w and represented by the operator $E_w\{\}$. This is estimated by the probability weighted sum of costs calculated from discrete realizations $w_{(k)}$, thus

$$f = \sum_k \{p_k [C(x, u, w_{(k)}) + F(y_{(k)})]\}$$

for weights p_k , where $y_{(k)} = T(x, u, w_{(k)})$.

This function allows inclusion of uncertain costs in the current stage; however, this can be simplified to

$$f = C(x, u) + \sum_k \{p_k F(y_{(k)})\}$$

to consider only deterministic current costs by choosing stages that are sufficiently short. *Foufoula-Georgiou and Kitanidis [1988]* also use this practical simplification.

C. PROBLEM DESCRIPTION

The sample problem models and optimizes control decisions for a conjunctive-use system. The problem contains no heuristic control rules that restrict the management of groundwater, and thus the optimal control solution has a greater set of feasible decisions from which to find the optimal. This allows us to observe the effect that the different capabilities and constraints of reservoir and aquifer storage have on control decisions, especially with respect to hedging of supplied water and allocation of stored water. In addition we can observe how the persistence of stream flows affect hedging and allocation.

The sample problem uses a generic conjunctive use model similar to that discussed by *Buras* [1963, 1972] and summarized by *Yakowitz* [1982]. It is based on the EBMUD system that supplies water to communities along the eastern edge of the central and southern San Francisco Bay. Figure II-1 illustrates the EBMUD system and conjunctive use model.

1. System Description

The EBMUD system serves the residential needs of approximately 1.2 million people, as well as the industrial, commercial, and institutional needs in the East Bay region of the San Francisco Bay area. Notable cities in the area include Oakland and Berkeley. About 95 percent of its water supply is from the Mokelumne River's 575-square mile watershed on the western slope of the Sierra Nevada Mountains.

Source

Supplies from stream flow are seasonal and uncertain, with flow into EBMUD's reservoir system at about 720 thousand acre-feet (TAF) annually (based on United States Geological Survey records since 1928 for the Mokelumne Hill Gage). Upstream of EBMUD's system there are small diversions of about 10 TAF annually by rural Amador and Calaveras Counties. Annual flow has varied from a high of 1,595 TAF in 1983 to a low of 130 TAF in 1977. Average monthly flows vary from a high of over 100 TAF in May to a low of about 30 TAF through the fall months. However, a typical year will generally have a month of high flow approaching 200 TAF and a month of low flow approaching 10 TAF; and, at times, the natural flow of the river has ceased.

Knowledge of monthly flow distributions, autocorrelation, and snow accumulation provides valuable information for improving predictions of flow. Intra-annual stream flows are highly autocorrelated. Also, much of the late spring and summer runoff is a result of melting snowpack, and measurements of snowpack are available throughout the winter. Therefore, consideration of the season and measurements of snowpack and prior stream flows all contribute to improvements in stream-flow prediction.

Existing System

EBMUD manages two reservoirs having a combined capacity of 641 thousand acre-feet (TAF) on the Mokelumne River. Up to 200 TAF of this storage is reserved for flood control under agreements with the U. S. Army Corps of Engineers, and an additional 21 TAF is dead storage. Besides water supply and flood control functions, the reservoirs also have a combined hydropower generating capacity of 39 megawatts.

An 82 mile aqueduct transports water to the service area by for use or storage in five terminal reservoirs. The terminal reservoirs provide an additional 155 TAF of storage capacity (with 17 TAF of dead storage). The aqueduct has a delivery capacity of 200 million gallons per day (MGD) by gravity flow or 325 MGD by pumping. This maximum delivery capacity also coincides with EBMUD's permits to divert up to 325 MGD (364 TAF per year) for use in its service district.

Proposed Aquifer Storage

Both the River and Aqueduct flow west from the Sierras and across the Central Valley of California before meeting the San Francisco Bay Delta. This area contains large, extensive aquifers.

Underlying this area, the fresh-water bearing Victor, Laguna, and Mehrten formations generally consist of thick Quaternary and late Tertiary sand and gravel deposits. Under these are brackish and saline marine sediments. Though locally confined conditions exist, the upper formations generally form an unconfined aquifer. Thicknesses of these formations total a few hundred feet in the east increasing to almost two thousand feet near the Delta. Well capacities have been reported in the range of 500 to 1500 gallons per minute (gpm), and studies indicate specific capacities in the range of 35 to 60 gpm/ft and transmissivities of 60,000 to 80,000 gal/day/ft [DWR, 1967].

Due to pumping for local agricultural and municipal needs, the water table is depressed by an average 50 to 100 feet below virgin levels. Currently the water table averages about 100 feet below land surface and is below sea level in many locations. This has resulted in salt-water intrusion near the Delta and has prompted local agencies to seek remedies that reduce over-pumping in the area. However, as a result of this over pumping, there is a great amount of available space for aquifer storage.

Additional aquifer storage opportunities exist in the EBMUD service district and other areas adjoining the District. Indeed, neighboring Alameda County Water District already has a local program for aquifer replenishment, though its efforts are aimed at restoring groundwater levels and preventing salt water intrusion rather than at managing the aquifer as a storage mechanism [Sax *et al.*, 1991]. Nevertheless, EBMUD has identified the lower Mokelumne region in the Central Valley as the best site based on storage capacity, well capacities, water quality, institutional constraints, and anticipated costs [EBMUD, 1992, Appendix D3].

2. Optimal Control Problem

The conjunctive-use model is a simplified representation of the EBMUD system with an added aquifer-storage component. The model does not solve specific management or planning problems for EBMUD, but it does illustrate some of the general characteristics of optimal control of a conjunctive-use system and points to differences between reservoir and aquifer storage.

System Model

I model the system with three state variables. These represent reservoir storage level (x_1), aquifer storage level (x_2), and prior month's stream flow (x_3). Control decisions specify the amount of water supplied (u_1) to meet demands and the change in groundwater storage (u_2) resulting from recharging and pumping the aquifer. These decisions are applied over stages that are one month long. During each stage, the state evolves from beginning values, x , to the ending values, y , based on control decisions and realized value of uncertain stream flow (w_1).

This effort is the first to illustrate the effect that the different capabilities and constraints of reservoir and aquifer storage have on hedging and allocation, without the use of heuristic rules controlling the use of groundwater. In

addition, this effort includes the effect of autocorrelation and illustrates the influence that autocorrelation can have on hedging and allocation.

Nevertheless, a practical and complete model of EBMUD's potential conjunctive-use system is more complex, with more state and decision variables, than presented here. A more complete model could require six or more state variables (Table 1). Here I use only three; I limit my current effort primarily to illustrate more easily a general conjunctive-use modeling and solution approach and to determine some general properties of conjunctive-use solutions.

Transfer Equation

Reservoir levels change linearly with stream flow and with control decisions that supply and allocate water. Groundwater level changes with the allocation decision that recharges or pumps. The transfer equation incorporates these effects as follows

$$y = T(x,u,w) = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix} x + \begin{bmatrix} -1 & -1 \\ 0 & 1 \\ 0 & 0 \end{bmatrix} u + \begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix} w .$$

Constraints

I incorporate fundamental differences between reservoir and aquifer storage by adding constraints on control decisions. As discussed above aquifers generally have much greater storage capacities than reservoirs but more limited response rates from filling and extraction.

We may be able to design reservoirs with storage capacities approaching that of many available aquifers; we may also be able to design pumping and recharge facilities that allow rapid changes in groundwater levels. However, a system with reservoir and aquifer storage can make effective use of both storage mechanisms without building facilities needed to remove these constraints. By efficiently combining the two storage mechanisms, we can achieve a high level of water supply reliability with much less expense.

Table 2 summarizes the constraints that I apply to the conjunctive-use problem. These constraints account for differences between reservoir and aquifer storage by specifying maximum reservoir level and maximum change in aquifer level during a one-month-long stage. In standard form, these constraints are:

(a) Upper bound of reservoir capacity: $[-1, -1]u \leq 400 - x_1 - w_1$

- | | |
|--|-----------------------------------|
| (b) Lower bound of reservoir capacity: | $[1, 1]\mathbf{u} \leq x_1 + w_1$ |
| (c) Lower bound of aquifer storage: | $[0, -1]\mathbf{u} \leq x_2$ |
| (d) Upper bound on recharge: | $[0, 1]\mathbf{u} \leq 10$ |
| (e) Upper bound on pumping: | $[0, -1]\mathbf{u} \leq 20$ |
| (f) Upper bound on demand: | $[1, 0]\mathbf{u} \leq 58.3$ |
| (g) Lower bound on demand: | $[-1, 0]\mathbf{u} \leq 0$ |

where numbers represent thousands of acre-feet (TAF).

Demand is assumed to be 700 TAF per year. This is close to the mean annual stream flow. Thus, demand in each stage is 58.3 TAF since control decisions are updated monthly.

The stochastic variable in these constraints presents a difficulty. An inequality constraint becomes an equality constraint when binding, and the presence of a stochastic variable in a constraint means that there is not a unique control solution that will satisfy the constraint for each possible realization of the stochastic variable.

I use two assumptions to address the difficulty of stochastic variables in the constraints. This first assumption is that demands can be met only by stored water and not by current stream flow. As a result, the supply and allocation decisions are more conservative than necessary, but this avoids the difficulty of specifying decisions that change throughout the month to adjust to changing stream flow. The second assumption is that the discrete realization of stream flow that first produces a binding constraint is the only value needed to specify a stochastic constraint. For example, stochastic constraint (b) is represented by a set of constraints resulting from each discrete realization of w_1 :

$$\mathbf{g}^T \mathbf{u} \leq x_1 + w_{1,(1)} \quad , \quad \mathbf{g}^T \mathbf{u} \leq x_1 + w_{1,(2)} \quad , \quad \mathbf{g}^T \mathbf{u} \leq x_1 + w_{1,(3)} \quad , \quad \dots$$

where $\mathbf{g} = [1, 1]^T$. If $w_{1,(1)}$ is the smallest discrete realization, only the first constraint can become an equality constraint; converting any other constraint in the set to an equality constraint is infeasible since this would violate the first constraint. Thus, $w_{1,(1)}$ is the only value needed to specify the stochastic

constraint (b). If stream flow, w_1 , can attain a minimum value of zero with non-trivial probability, it is reasonable to replace stochastic constraint (b) with

$$(h) \quad [1, 1]u \leq x_1 .$$

Control decisions will then be feasible for any realized value of stream flow.

A problem with these two assumptions is that control decisions can be excessively constrained. It is even possible that the constraints result in no feasible decisions. To prevent this, we can choose a decision period that is sufficiently short, and we can limit selection of discrete values $w_{(k)}$ to reasonable values that are contained in a bounded set. In the conjunctive-use problem presented here, any non-negative value of $w_{(k)}$ is possible, though the probability of extremely large values becomes insignificant and does not affect the optimal solution. Because the objective function includes only demand for water (and not, for example, flood control), releases in excess of demand required to meet the reservoir capacity constraint (a) have no effect on the objective and optimal controls. If a release is required, the reservoir is full and thus the amount released does not affect future decisions or costs. Thus, for this problem, I drop constraint (a) and set $y_1 = 400$ whenever the transition equation results in $y_1 > 400$.

An alternate approach to handling stochastic constraints that I do not use here is to add stochastic slack variables s to the transition equation. This results in the transition equation

$$y = T(x, u, w) = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix} x + \begin{bmatrix} -1 & -1 \\ 0 & 1 \\ 0 & 0 \end{bmatrix} u + \begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix} w + \begin{bmatrix} -1 \\ 0 \\ 0 \end{bmatrix} s .$$

The values of slack variables depend on control decisions and realized stream flow, thus $s(u, w)$. Though more complicated, the solution algorithm presented in Chapter IV could be expanded to include these variables.

Objective Function

The objective function is the recursively derived function

$$f = C(u) + \sum_k \{p_k F(y_{(k)})\}$$

where

$$C(u) = (5 - 4u_1/58.3)(58.3 - u_1) + 0.0001u_2^2$$

and

$$F_{(N)}(y) \equiv 0 .$$

58.3 is the monthly water demand. This objective function only incorporates the cost of water rationing. It does not include other costs of flooding or groundwater pumping, or benefits of hydropower generation. It does, however, include a small penalty for groundwater pumping and recharge to make the objective strictly convex with respect to these decisions. This prevents unnecessary pumping and recharge.

Though not based on a practical assessment of rationing costs, the objective function contains many characteristics of a practical function. Significant characteristics of any practical function are increasing total and marginal costs as rationing increases (supply decreases). This is included in the objective function.

This function assumes that an unlimited, fixed-price water supply alternative does not exist. As a result, the function is strictly convex with respect to rationing decisions, and optimal decisions prefer long periods of moderate rationing to short periods of severe rationing. Was I to use a linear function representing the presence of a fixed-price water supply alternative, rationing would not occur beyond a level that produces marginal costs equal to the fixed price. In practical problems, this condition rarely exists. However, even if I was to consider this possibility, the objective function would still be convex, but not strictly convex. Concave functions are rare; optimal control with a concave objective function can produce an unlikely all-or-nothing approach to allocating supplies.

Use of the above objective function is also justified since it incorporates rationing costs, the most significant component of operating cost for many systems. Most other costs are less significant when compared to rationing costs. For a conjunctive-use system with wells pumping at 1000 gpm and Central Valley aquifer conditions discussed above in Section C.1, pumping costs are about \$10 to \$20 per acre-foot based on just the cost of electricity [Georgakakos and Vlatza, 1991]. For the EBMUD system, the historical value of hydropower generated by the EBMUD system is about \$10 per acre-foot [EBMUD, 1992, Vol. 1]. In contrast, when water is scarce, the marginal value of municipal water supplies can be much greater. California sold water from its 1992 Water Bank for \$175 an acre-foot [McClurg, 1992b]. The California communities of Santa Barbara, Goleta and Montecito, and others have developed desalination plants to

provide backup water supplies that cost \$1500 to \$2000 an acre-foot. Even when water is not scarce, wholesale costs south of the Delta range from \$44 in San Joaquin Valley to \$237 in southern California [McClurg, 1992a].

3. Stochastic Model of Stream-Flow Time Series

Figure II-2 displays the monthly stream-flow hydrograph of the Mokelumne River. I model the probability distributions of monthly stream flows from this time series by using a reasonable three-parameter model. The model also incorporates an additional parameter for each month to account for stream-flow autocorrelation with the prior month. For a practical model of the EBMUD system, a stream-flow model could also include measurements of snowpack accumulation during the winter and spring; the approach used to incorporate this additional information would be similar to that used to incorporate autocorrelation.

Model of Distribution

The probability distribution for each months stream flows are modeled by a seasonal 3-parameter, log-normal model. Chapter III discusses how seasonality is incorporated in the optimal control problem. Parameter χ_τ shifts stream-flow values to fit more closely a log-normal distribution modeled by log-mean μ_τ and log-standard deviation σ_τ , thus

$$w_\tau = \chi_\tau + \exp(\varepsilon\sigma_\tau + \mu_\tau)$$

for each month τ . Stochastic variable ε has a standard-normal distribution. These parameters are estimated from the historical data as discussed below. Examples of probability quantiles for monthly flow are displayed in Figure II-3.

Model of Autocorrelation

Autocorrelation is first-order auto-regressive (Figure II-4). Correlation coefficients are calculated after stream-flows have been transformed because this produces coefficients that are more significant than those calculated prior to transformation. Thus, the dependence of w_τ on $w_{\tau-1}$ (where $w_{\tau-1} = x_{3,\tau}$) is modeled by the relationship

$$w_\tau = \mu_\tau + \rho_\tau \frac{\sigma_\tau}{\sigma_{\tau-1}} (w_{\tau-1} - \mu_{\tau-1}) + \varepsilon\sigma_\tau \sqrt{1 - \rho_\tau^2}$$

where $\omega_\tau = \ln(w_\tau - \chi_\tau)$ and $\omega_{\tau-1} = \ln(w_{\tau-1} - \chi_{\tau-1})$. Parameter ρ_τ is the sample correlation coefficient for month τ .

Parameter Estimation

The correlation coefficient is estimated from historical data by

$$\rho_\tau \equiv \frac{\text{Cov}(\omega_\tau, \omega_{\tau-1})}{\sqrt{\text{Var}(\omega_\tau) \text{Var}(\omega_{\tau-1})}}.$$

Stedinger [1981] discusses a variety of additional definitions for the correlation coefficient that also may be appropriate. Parameter values μ_τ and σ_τ are the unbiased sample estimates from the transformed historical data. The third "shift" parameter, χ_τ , of the 3-parameter log-normal distribution is estimated by

$$\chi_\tau \equiv \frac{w_i(q) w_i(1-q) - w_i(0.5)^2}{w_i(q) + w_i(1-q) - 2w_i(0.5)}$$

as recommended by *Stedinger* [1980], where parameter $w_i(q)$ is the sample value for the q 'th quantile. *Stedinger* recommends choosing $w_i(q)$ and $w_i(1-q)$ to be the largest and smallest values observed for month τ . Table 3 summarizes the parameters calculated for each month.

Considerations

An advantage of the three-parameter log-normal model is that, for the example problem, $\partial w / \partial x$ becomes

$$\left[\begin{array}{ccc} 0, & 0, & \frac{w_1 - \chi_\tau}{x_3 - \chi_{\tau-1}} \rho_\tau \frac{\sigma_\tau}{\sigma_{\tau-1}} \end{array} \right].$$

The three-parameter log-normal model also presents some possible complications. These include

1. Values of a historical data set that result in $w_\tau - \chi_\tau < \delta$ where δ is a small value: Using the recommendation of *Stedinger* [1980] ensures that the estimate of χ_τ will be smaller than the smallest observed value of w_τ . However, the discretized state space may include values, $w_{k,\tau}$, that could violate this condition. Thus, I do not allow values of χ_τ greater than -1. The Mokelumne River flow series generally is slightly less skewed than a fitted log-normal distribution; as a result, the samples χ_τ are negative for all months but July (Table 3).

2. Historical data that are not positively skewed: A three-parameter log-normal model cannot produce a negative skew; however, without constraints the above parameter estimates will try to fit a negative skew. To prevent infeasible parameters, a probability distribution model must satisfy the restriction $w_i(q) + w_i(1-q) - 2w_i(0.5) > 0$ [Stedinger, 1980]. As skew approaches zero, χ_τ approaches $(-\infty)$. For the Mokelumne River flow series, all monthly distributions have a positive skew except for September and October. χ_τ for these months is set to -600 to more closely approximate a skew of zero (Table 3).
3. Synthetic stream-flow values that are negative: When this occurs, the generated flow value and $\partial w_1 / \partial x_3$ are set to zero. Since $\chi_\tau \leq -1$, generated flow values may be shifted to less than zero.
4. Synthetic and real stream-flow values that are greater than the largest value of the discretized range of the distribution: To discretize the state-space of the system, a bounded range of flows must be specified; however, log-normal distributions are unbounded. Fortunately, optimal decisions applied to water management problems are not sensitive to sufficiently large stream-flows for which the probability approaches zero. This is especially true for problems with only a water supply objective—optimal decisions do not change when stream flows are greater than the sum of storage capacity and demand since the system cannot make use of flows beyond this sum. Thus, when generated flows greater than this sum occur, I reduce their value to this sum and can set gradients of the objective function $\nabla_x f$ to zero.




Fig. II-1. EBMUD system and modeled conjunctive-use system.

Table 1. State Variables Needed by Models

	Model of Real System	Current Model
Facilities:	Reservoirs (2-3) Groundwater Storage (min 2)	Reservoir (1) Groundwater Storage (1)
Stochastic Inputs:	Prior Stream Flow (1) Snow-Pack Measurements (1)	Prior Stream Flow (1)
Total	min 6	3

Table 2. Constraints on State and Decision Variables
(1000's of acre-feet)

	Minimum	Maximum
<i>Surface Reservoir</i>		
Storage	0.0	400.0
Change per month	unlimited	unlimited
<i>Sub-Surface Reservoir</i>		
Storage	0.0	unlimited
Change per month	-20.0	10.0
<i>Release for supply</i>		
Supply per month	0.0	58.3




Fig. II-2. Stream-flow hydrograph of the Mokelumne River.




Fig. II-3. Quantiles of the fitted 3-parameter log-normal distribution. Displayed quantiles are stream-flow values resulting from the back-transformation of the first and second standard deviates of the normalized residuals, where

$$w_{\tau}(\epsilon) = \chi_{\tau} + \exp(\mu_{\tau} + \epsilon\sigma_{\tau})$$

for each month, τ , and deviate:

$$\epsilon = -2, -1, 0, 1, 2$$

and resulting quantiles:

$$P\{w_{\tau} \leq w_{\tau}(\epsilon)\} = 0.054, 0.242, 0.5, 0.758, 0.946$$

Autocorrelation and Partial Autocorrelation of Monthly Stream Flows




Fig. II-4. Lag-one and lag-two autocorrelation coefficients, $\rho_{\tau}(1)$ and $\rho_{\tau}(2)$, and lag-two partial autocorrelation coefficient, $\pi_{\tau}(2)$, for each month τ . Partial autocorrelation is the correlation that remains after accounting for shorter lags. Thus $\pi_{\tau}(1) = \rho_{\tau}(1)$, and $\pi_{\tau}(2) = [\rho_{\tau}(2) - \rho_{\tau}(1)\rho_{\tau-1}(1)]/[1 - \rho_{\tau}(1)\rho_{\tau-1}(1)]$.

Table 3. Stream-flow Parameters

Parameters that transform stream flows, w , to a standard-normal distribution using the 3-parameter lognormal model:

$$\varepsilon = [\ln(w-\chi) - \mu]/\sigma$$

Month	shift parameter χ	normalized mean μ	std dev. σ	lag-1 a.c. ρ
Oct.	-600 *	6.448	.018	.928
Nov.	-6.806	3.617	.527	.812
Dec.	-1.106	3.627	.736	.776
Jan.	-3.512	3.824	.698	.807
Feb.	-3.029	3.875	.723	.763
Mar.	-10.101	4.227	.573	.821
Apr.	-24.363	4.569	.477	.824
May	-90.159	5.271	.377	.896
Jun.	-35.207	4.824	.537	.821
Jul.	-1 **	3.573	.612	.700
Aug.	-89.184	4.804	.094	.917
Sep.	-600 *	6.447	.016	.935

* distribution has negative skew, χ value at -600

** value adjusted to -1

D. SOLVING THE CONJUNCTIVE-USE PROBLEM

1. Method

The optimization method used to make iterative improvements on the solution is a well-established optimization technique known as the method of successive approximations [Esmail-Beik and Yu, 1984]. This method solves the recursive relationship

$$F_{(t-1)}(y_{(t-1)}) = f^*_{(t)}(x_{(t)})$$

where
$$f^*_{(t)}(x) = \min_u \{ C(x,u) + \sum_k \{ p_k F_{(t)}(y) \} \}$$

backwards in time from a specified expected cost function $F_{(N)}(y)$. This cost function specifies expected costs beyond the end of a planning horizon, N stages long. Note that $y_{t-1} \equiv x_t$.

If we have sufficient experience with solving a particular class of problems and know the solutions present some simple characteristic functional form, we can describe F as a function of y without using the piece-wise interpolation method used here. The problem solving would then consist of a few fitted parameters. Often, however, we cannot specify a simple characteristic functional form unless we are willing to make some very broad assumptions.

Instead, we can specify the value of F at a number of discrete states $y_{(i)}$ sufficient to allow us to estimate the continuous function by interpolation. Using this piece-wise continuous expected cost function, we can solve the recursive relationship presented above by iteratively improving control decisions until the optimum is identified (Figure II-5).

2. Interpretation

Suppose that we are extremely short-sighted in our management of a system, and we chose the goal of minimizing cost for just the current stage. By ignoring future costs, we are—in effect—saying that our expected cost function is zero (or constant) for all final states of the system: $F_{(N)}(y) \equiv 0$. Optimal decisions u^* are those that give

$$f^*_{(N)}(x) = \min_u \{ C(x,u) \} .$$

If we determine the optimal decisions and cost for discrete initial states $\mathbf{x}_{(i)}$, we can specify a piece-wise continuous optimum control function $\mathbf{u}^*(\mathbf{x})$ and optimal cost function $f^*_{(N)}(\mathbf{x})$ for all states $\mathbf{x}$.

Now suppose that we want to be a little less short sighted, and we decide to minimize the total costs incurred over both the current stage and a following stage. In the first stage of this two stage problem, optimal decisions are those that give

$$f^*_{(N-1)}(\mathbf{x}) = \min_{\mathbf{u}} \{ C(\mathbf{x}, \mathbf{u}) + \sum_k \{ p_k F_{(N-1)}(\mathbf{y}_{(k)}) \} \} .$$

Under this goal, it is, for example, no longer efficient for the manager of a water supply system to drain the reservoir to meet current demands if that decision results in a severe shortage in the next stage. To solve this problem, we now need an estimate of $F_{(N-1)}(\mathbf{y})$ representing the expected costs of one following stage. Assuming that optimal decisions will also be applied to this stage, this expected cost function is simply the optimal cost function of the previous single period solution; i.e.,

$$F_{(N-1)}(\mathbf{y}_{N-1}) = f^*_{(N)}(\mathbf{x}_N)$$

where $\mathbf{y}_{N-1} \equiv \mathbf{x}_N$. If we again determine the optimal decisions and cost for discrete initial states $\mathbf{x}_{(i)}$ of the system, we can specify a piece-wise continuous optimum control function $\mathbf{u}^*(\mathbf{x})$ and associated optimal cost function $f^*_{(N-1)}(\mathbf{x})$ for a two-stage model.

We can repeat this process recursively replacing $F_{(N-1)}$ by $F_{(N-2)}$ by $F_{(N-3)}$ and so on, until we incorporate the costs of all N stages representing the planning horizon. Each recursion steps further back in time and away from the zero-valued expected cost function; and the resulting optimal cost function represents expected costs for a planning horizon that is one stage longer.

3. Simple Illustration

The method is clearly illustrated by the control of simple water supply system. This system consists of a single reservoir with capacity $x_{\max}$ and an uncorrelated stream-flow source. The state of the system is described by the reservoir level x , and the only control decision u is how much to release. The transition equation is

$$y = x - u + w$$

where w is stream flow and y is the resulting reservoir level.

Constraints are $x_{\max} \geq y \geq 0$ and $u \geq 0$. In standard form these become

$$u \leq x + w$$

$$-u \leq x_{\max} - x - w$$

and

$$-u \leq 0 .$$

Assuming that planned releases can only use previously stored water, we can ignore stream flow in the first constraint. Also, for an objective considering only water supply, the second constraint can be dropped since spills from a full reservoir have no affect on optimal control. If the transition function produces a final reservoir level greater than the maximum, set $y = x_{\max}$.

When a release is less than demand, the system incurs a cost $C(u)$.

Assuming $F_{(N)}(y)$ is zero for any ending reservoir level, there are no costs beyond the end of the planning horizon. The only costs that affect control decisions are those accrued between the current and last stages.

For example, if there is only one remaining stage, the efficient release $u^*(x)$ is that which achieves

$$f_{(N)}^*(x) = \min_u \{C(u)\} .$$

Because $F_{(N)}(y)$ is zero, there is no benefit from saving water in the reservoir for future use. Thus there is no incentive to hedge. The efficient decision, u^* , is to release as much as possible to meet demand, subject to the constraint $u \leq x$ (Figure II-6).

If, on the other hand, there are two remaining stages, the efficient release becomes that which achieves

$$f_{(N-1)}^*(x) = \min_u \{C(u) + F_{(N-1)}(y)\}$$

subject to the constraints and the transition equation. Since $F_{(N-1)} = f_{(N)}^*$,

$$f_{(N-1)}^*(x) = \min_{u'} \{C(u') + \min_u \{C(u)\}\} .$$

There now is an incentive to hedge in the current period if the cost of hedging is less than the benefit of the extra water in the last period. As expected costs

incorporate more stages and a longer planning horizon, hedging becomes more common (Figure II-7).

4. Stochastic Versus Deterministic

The solution of the above problem produces an expected cost function, $F(y)$, for every stage of a planning horizon. We can think of these expected cost functions as the sum of a series of current costs, C , resulting from a control decision applied to each stage remaining until the end of the planning horizon. If all inputs to the system are certain, we can specify, in advance, optimal decisions and states of the system by a single control schedule. The cost function is a sum of known costs

$$F_{(0)} = \sum_{t=1}^N \{C(x_t, u_t)\} .$$

However, if there are uncertain inputs, a single control schedule cannot specify the optimal control of a system and the cost function is a sum of expected, vice known, costs.

If there are uncertain inputs w , one cannot know the state of following stages until the beginning of those stages. Thus we cannot determine optimal controls and costs until the beginning of a stage. A control schedule applied to a system depends on unknown future states of the system resulting from realized values of w . Because the state at any particular stage is unknown, a solution must specify optimal controls and costs for any possible state.

This is the major difficulty in solving stochastic dynamic programming (SDP) problems. Because we must specify a solution for every possible state of a system, we potentially have a huge number of calculations to perform and a huge amount of data to store. If a solution can be specified by a simple function, the parameters of the function specify the solution. Generally, however, we cannot specify a solution by a simple function. Instead, we must approximate the solution by calculating values for many discrete states—this is the Discrete Dynamic Programming (DDP) method. To maintain a particular level of solution accuracy, the number of discrete states increases exponentially with each additional variable added to the state vector. Problems with more than two or three state variables generally present a great computational challenge. This can be addressed partly by improved optimization methods, such as discussed in Chapter IV. Also we can reduce the challenge by understanding better the

characteristics of a solution, thus we may be able to apply simplified models and functions that reduce the number of discrete states.

Conjunctive-use problems necessarily start with at least three state variables, and this is when the model that is generally unrealistically simple. A conjunctive-use model requires, at a minimum, state variables to represent levels in reservoir and aquifer storage, and stream flows. Generally, we can model the state of a reservoir by a single variable for storage, though each reservoir needs its own state variable unless we aggregate their characteristics. On the other hand, it is not clear that the state of aquifer storage could ever be realistically modeled by a single state variable since pumping costs depend on the distributed characteristics of the aquifer. Stream flows are usually autocorrelated for the length of stages practical for real-time system control, and state variables must be added to incorporate information from prior flows and from other sources such as snow-pack measurements. For example, Table 1 outlines the state variables that may be required to model the EBMUD system. Because of the number of state variables required to model realistic conjunctive-use problems, it is unlikely that such problems can be solved by traditional SDP methods.

I use Gradient Dynamic Programming (GDP) of *Foufoula-Georgiou and Kitanidis* [1988] to solve the problem. GDP is a new method that allows us to solve problems with a greater number of state variables. Chapter IV presents the method. The algorithm that I present is an extension of the original method to include autocorrelation.

5. Convergence

The effects of current decisions are preserved and transmitted to later stages through the state variables y that result from the decisions and inputs applied to the initial state x . $F(y)$ is the sum of all expected costs starting from state y and applying optimal decisions in remaining stages. In a problem with stochastic inputs, we can determine $F(y)$ by recursively solving problems with increasing numbers of stages until $F(y)$ represents expected costs for an entire planning horizon.

The cost function $F(y)$ influences control decisions by providing the relative costs of different decisions. The absolute value of $F(y)$ does not influence control decisions, thus we can re-define $F(y)$ as

$$F_{(t-1)} = f_{(t)}^* - c_{(t)}$$

without affecting control decisions. Constant $c_{(t)}$ standardizes the cost function, allowing us to more easily view changes with time, $t = N, N-1, N-2, \dots 1$, and avoiding potentially large-valued functions. Also, with an appropriate definition, we can use $c_{(t)}$ as an indication of the overall change in the cost function, such as the extra cost expected for an additional stage of operation.

Stedinger et al. [1984] suggest setting $c_{(t)}$ to the average gain in costs incurred during stage t . Another nearly equivalent approach is to set $c_{(t)}$ to the mean or other measure of the central tendency of $f_{(t)}^*$. A third approach is to set $c_{(t)}$ to the smallest discrete value of $f_{(t)}^*$. With this approach $F_{(t-1)}$ is a strictly positive function; though, since $c_{(t)}$ is based on a single discrete value instead of an average of all discrete value, numerical errors can be more significant.

If, after N periods, the assumed cost function $F_{(N)}(y)$ has a negligible effect on current decisions, $F(y)$, $u^*(x)$, and c should converge (Figure II-8). This will often be the case for water projects where the planning horizon may be tens of years. In these cases, we can solve the problem with fewer recursions by performing only as many as required to obtain convergence.

When convergence is achieved, the value of c is a measure of stage to stage costs. This parameter can be used to compare performance of different system configurations. For example, this parameter can be used to compare the benefit of adding aquifer storage or additional reservoir storage to a system.

Also, the current state of the system has only a transient effect on the expected cost of the system; if we look far enough into the future, we expect costs for each additional period of operation to be the same regardless of the current state of the system. For example, though the system may currently experience drought or flood, the expected cost of future operations will approach the long-term average. The number of stages required to achieve convergence indicates the system's memory of control decisions; thus, this could provide some insight into differences between systems. For example, limits on pumping and recharge rates will generally result in a slower convergence. This is related to groundwater's relative weakness in meeting acute short-term water supply needs but its relative strength in meeting moderate long term water supply needs.

In this thesis, c is the smallest discrete value of $f_{(t)}^*$. I use c only as a standardizing constant to make it more easy to view changes in $F(y)$; we prefer the advantage of a strictly positive $F(y)$ over a better measure of change.

Recursive Development of a Control Solution




Fig. II-5. Real time of decisions progresses to the right while development of solution recurses to the left. For the simple water supply system of figure 1, the sequence of optimum transitions are those that might be expected for a dry-wet-dry sequence of stream flows. The arrows do not connect because we develop a cost-to-go function that is interpolated between the discrete values. Thus it is not necessary to require that transitions end on the discrete states used to span the state variables.

Optimal Decisions Applied to the First Recursion


Optimal decisions minimize current and future costs:

$$\min\{ f(x) \} = \min\{ C(u) + F(y) \}.$$

The first recursion is the last stage, so $F(y)$ is zero for all ending states y :

$$\min\{ f(x) \} = \min\{ C(u) \}.$$

For each discrete reservoir level x_i , we determine a release decision, u_i , that minimizes current cost $C(u_i)$:



This results in an expected total cost $f(x_i)$ and change in reservoir level:





Fig. II-6. Illustration of optimal release decisions for the first recursion (last stage) and the resulting expected costs and changes in reservoir level.

Optimal Decisions Applied to the Second Recursion

In the second recursion, the expected cost function, $F(y)$, is no longer zero. It is now equal to the expected total cost function of the first recursion. Thus,

$$\min\{f(x)\} = \min\{C(u) + F(y)\}$$

For each discrete reservoir level x_i , we again determine a release decision, u_i , that minimizes current cost $C(u_i)$:



This again results in an expected total cost $f(x_i)$ and change in reservoir level, but now release decisions may hedge to avoid more severe rationing in the following stage:




Fig. II-7. Illustration of optimal release decisions that incorporate hedging when expected future costs are not zero.

Solution Convergence




Fig. II-8. Increases in the expected costs with the number of remaining stages should converge to a constant rate, regardless of the initial state of the system. Initial states x_1 , x_2 , and x_3 illustrate this, where x_1 is more costly than x_2 which is more costly than x_3 .

E. RESULTS

The solution of the example problem demonstrates four important observations about the optimal control of conjunctive-use systems: (1) efficient control requires effective use of the capabilities of both reservoir and aquifer storage; (2) autocorrelation significantly affects optimal control; (3) analytic optimization of conjunctive-use systems is important and perhaps necessary for achieving optimal control; and (4) the solution of the conjunctive-use management problem has characteristics that may simplify future studies.

1. Efficient Allocation

The solution demonstrates that optimal management of conjunctive-use systems takes advantage of the strengths of both mechanisms, achieving greater system efficiency and reliability than alternate systems using either mechanism in isolation. Specifically: (1) optimal management results in lower rationing costs than achievable by heuristic management (such as using aquifer storage only as a back-up); and (2) optimal allocation results in a preference for aquifer storage over reservoir storage as the risk of near-term rationing declines.

We can most easily observe these two results by viewing contours of expected rationing cost. Figure II-9 displays an example set of contours plotted against reservoir and aquifer storage levels. These contours indicate higher costs for lower storage levels. Also, the surface defined by these contours is convex.

Efficiency

Given a total amount of stored water, we can identify a minimum cost allocation between reservoir and aquifer storage. A general principle of economics is that decisions are made at the margin: if the benefit of one additional unit of groundwater is greater than the benefit of one additional unit of water in the surface reservoir, we would prefer to recharge groundwater over filling the reservoir. For example, the allocation of 300,000 acre-feet of water is constrained by a linear function A--A' in Figure II-9. Since the cost curves are convex, we can easily identify a minimum cost allocation. For 300,000 acre-feet of water, the minimum cost allocation is at point B.

Given expected costs as in Figure II-9, we can identify the lowest cost allocation using both reservoir and aquifer storage. However, we cannot ensure that this will be achieved by control decisions because of uncertain stream flows.

The precise result of control decisions is determined by values of stream flow realized in the future: actual flows that are higher than expected result in over allocation to the reservoir, actual flows that are lower than expected result in over allocation to the aquifer. Also, we cannot ensure that we will achieve the lowest cost allocation because of constraints on control decisions that may limit our ability to reach this allocation.

Though we cannot ensure that the lowest cost allocation will be achieved by control decisions, we can apply control decisions that attempt to reach this allocation. The optimization approach presented in this work attempts to reach this allocation, subject to the constraints and operational costs that these decisions entail. In contrast, control decisions determined heuristically do not recognize the optimal allocation and thus do not attempt to reach it. Thus, heuristic control rules generally are less efficient than rules constructed from the solution of an optimization problem.

Allocation

Given the ability to identify the least cost allocation of water, we can determine the optimum allocation for any total amount of stored water. Figure II-10 reorients the display of contours in Figure II-9 to display more clearly the least cost allocation of any total amount of stored water. The least cost allocations occur at the trough in the cost contours as indicated by a line in Figure II-10. We can observe the least cost allocation favors storing most water as groundwater, except when storage levels are very low.

These contours specify costs expected in November when stream flow in October is above average. The least cost allocations for these contours favor groundwater because (1) the likelihood of near-term rationing is low in November at the beginning of the winter rainy season and (2) the likelihood that future stream flows will spill (and be lost) from a full reservoir is high if more water is stored in the surface reservoir. Since the model assumes that current demand for water can be met only from storage and not from current stream flows, it is necessary to store a small amount of water in the reservoir.

2. Effect of Autocorrelation

When stream flow time series are autocorrelated, a below-average monthly stream flow decreases the expected value of later stream flows.

Expected costs increase since lower expected stream flows increase the likelihood of rationing.

Figures II-11 displays expected cost contours for November when stream flow in October is below average. Compared to Figure II-10, expected costs are greater for all combinations of storage levels since prior stream flow is 20,000 vice 60,000 acre-feet for October.

Autocorrelation influences the expected costs and consequently the optimal decisions that allocate water to users and between storage mechanism. Optimal control of the conjunctive-use system changes decisions that allocate and ration water for users and that allocate water between reservoir and aquifer storage.

Allocation to Users

System operators must continuously update control decisions based on the realized and expected values of inputs. With each decision, system operators make control decisions that balance current performance against uncertain future performance. Under these conditions, rational operators will hedge their control decisions to improve future performance.

As expected costs increase with decreasing stream flow, rational operators have a greater incentive to hedge by rationing the water supplied to users. This rationing reduces the possibility of more severe rationing in the future.

Figure II-12 is a surface plot of the release decision for various storage levels. The figure shows that less water is released as levels in both reservoirs decrease. With less water in storage, the likelihood of future shortages increases, increasing the expected costs. Thus, the surface plot of Figure II-12 has a shape that is roughly an inverse of that which would result from the expected costs of Figure II-11.

Figure II-12 displays the release decisions for November when stream flow in October is 20,000 acre-feet. For other months and prior stream flow, releases generally decrease with increases in expected cost, and vice-versa.

Allocation Between Reservoir and Aquifer Storage

Besides changing decisions that ration water supplies, stream-flow autocorrelation also influences optimal decisions that allocate water between the dissimilar mechanisms of reservoir and aquifer storage. Autocorrelation increases the time that prior stream flows influence future flows; however,

autocorrelation also decreases the severity of droughts due to the effective averaging of numerous realizations of stream flow. Thus, the degree of autocorrelation in a stream-flow time series significantly affects the efficient allocation of stored water between reservoir and aquifer storage.

As we saw above, Figures II-10 and II-11 display the expected cost contours for November when prior stream flows are above average and below average. We can see how the minimum cost allocations—indicated by lines connecting the minimum cost points of the contours—change with prior stream flow. As prior stream flows decrease, the likelihood of near-term water shortages increases. Comparing Figure II-11 to Figure II-10, we can see a reduced preference for aquifer storage due to the higher risk of near-term rationing. This is indicated both by the shift in the minimum cost allocation and by the flatter cost contours. As risk of near-term shortage increases, it is less likely that the reservoir capacity constraint will be binding on decisions, and the minimum expected cost is not sensitive to allocation. When the risk of near-term shortage is high, the only inefficient allocations are those that make exclusive use of aquifer storage: groundwater can meet only about a third of current demand due to the pumping rate constraint.

The minimum cost allocations of Figures II-10 and II-11 control the groundwater allocation decisions of Figures II-13a and II-13b. Figure II-13a displays recharge (positive values) and pumping (negative values) decisions for November when stream flow in October is above average. We can see that optimal control prefers groundwater recharge for all storage levels except when the reservoir is nearly empty. Comparing Figure II-13a to Figure II-10, we can see that when expected costs of rationing are low, we should store water in the aquifer to meet long-term needs. Conversely, Figure II-13b displays recharge and pumping decisions for November when stream flow in October is below average. We can see that optimal control prefers groundwater recharge less than in Figure II-13a. The differences in allocation decisions between Figures II-13a and II-13b reflect the differences in expected costs between Figures II-10 and II-11.

As a final note, Figures II-13a and II-13b display a dip in the recharge decisions as groundwater storage becomes large. This dip may be a result of the minor pumping and recharge costs applied to the objective function. The dip may also be an artifact of the optimization method or an error in the computer code used to obtain the solution: I have observed that some of the other control

surfaces have anomalous oscillations. I hope that subsequent work will identify the cause of these oscillations.

3. Optimization and Simulation

Burges and Maknoon [1975] note that "An important direction for research effort would be to examine the general nature of 'response surfaces' for conjunctive use problems. If they turn out to be relatively flat then it may be possible to use the simulation approach rather than direct optimization procedures. This would overcome many of the problems of dimensionality which limit the use of optimization models."

As we can see from the above results, the "response surfaces" for conjunctive use problems are sometimes flat and sometimes not. When the risk of water shortages is high in the example problem, the expected cost contours are flat and not sensitive to allocation between reservoir and aquifer storage. However, the example problem does not consider the costs of groundwater pumping or recharge. Including these costs will tip the cost contours to a greater preference for reservoir storage, especially when the risk of shortages is high.

Because of the different characteristics of reservoir and aquifer storage, we can conclude that the "response surfaces" for conjunctive-use problems are generally not "flat"—that is, the values of the objective and optimal control decisions change with the state of the system. Thus, optimization of conjunctive-use problems can significantly improve the control of conjunctive-use systems. Under these conditions, efficient heuristic control of conjunctive-use systems is unlikely. However, it may be possible to specify simple functions that capture characteristics of the response surfaces, thus reducing the number of discrete states that must be considered.

4. Possible Simplifications

One possible simplification that could be applied to future conjunctive-use problem solutions is the use of a global function as advocated by *Buras* [1972]. We can observe that the expected costs and control decisions for this stochastic problem are quite smooth. This is largely a result of the stochastic inputs which result in averaging over the possible input realizations. Because of this smoothness, it may be possible to describe the cost and control functions by a global polynomial or exponential function. Thus we could avoid some

discretization of the state space and solve similar problems with less computational effort.

Another possible simplification is the use of more constrained operating policies that simplify control of aquifer storage. *Bogle and O'Sullivan* [1979] observe that optimal operating policies derived by dynamic programming have a "common sense" form that results in a release that is "minimum for storage below some level and maximum for storage above some other level." Though the results presented here do not support such a simplification, there may be situations in which an operating policy presents a consistent form that can be used.

Expected Cost of Water Rationing Cost Versus Current Storage




Fig. II-9. Contour plot of expected costs (in arbitrary units) versus amount of stored water in reservoir and in aquifer.

The dashed line between A and A' indicates possible allocations of 300,000 acre-feet of total stored water. Point A represents allocation of all 300,000 acre-feet to reservoir storage with an expected cost of 520. Similarly, point A' represents allocation of all 300,000 acre-feet to aquifer storage with an expected cost of 460. Expected cost can be reduced to 370 at point B by allocating water to both reservoir and aquifer storage.

Minimum Cost Allocation of Stored Water
 Expected Cost Versus Total Amount and Allocation of Stored Water
 (prior month's stream flow above average)




Fig. II-10. Contour plot of expected costs versus total amount and allocation of stored water. Amount of water stored in reservoir and aquifer can be measured diagonally from zero. The results are for November in the first year of a four-year planning horizon when October's stream flow was above average (60,000 acre-feet).

The stippled line indicates the allocation of stored water that produces the lowest expected cost from rationing. Point B marks the optimum allocation for 300,000 acre-feet determined in Figure II-9.

Minimum Cost Allocation of Stored Water
Expected Cost Versus Total Amount and Allocation of Stored Water
 (prior month's stream flow below average)




Fig. II-11. Contour plot of expected costs versus total amount and allocation of stored water. The results are for November in the first year of a four-year planning horizon when October's stream flow was below average (20,000 acre-feet).

Supply Decision

Water Supplied Versus Current Reservoir and Aquifer Storage Levels




Fig. II-12. Surface plot of water supplied to meet a demand of 58.3 per month. Decisions to supply less than this amount indicate rationing. Results are for November in the first year of a four-year planning horizon when October's stream flow is 20. Amounts are in 1000's acre-feet.

Aquifer Storage Decision
Amount of Water Recharged (+) or pumped (-)
Versus Current Reservoir and Aquifer Storage Levels




Fig. II-13a. Surface plot of change in aquifer storage level. Positive values indicate recharge, with a maximum of 10 per month. Negative values indicate pumping, with a minimum of 20 per month (pumping cannot exceed 20,000 acre-feet/month). Results are for November in the first year of a four-year planning horizon when October's stream flow is above average (60,000 acre-feet). Amounts are in 1000's acre-feet.




Fig. II-13b. Surface plot of change in aquifer storage level. Results are for November when October's stream flow is below average (20,000 acre-feet). Amounts are in 1000's acre-feet.

F. CONCLUSIONS

I have illustrated the optimal control of a conjunctive-use problem with uncertain and autocorrelated inputs. From the results, we can observe that differences in the characteristics of reservoir and aquifer storage affect optimal control. This is especially true in the presence of inputs that are autocorrelated.

Results from an analysis, such as presented here, can provide practical information to planners and managers of water supply systems. Planners can anticipate the best choice of mechanisms to enhance the capabilities of a water supply system based on needs: reservoir storage likely will be a better choice for short term or seasonal needs; aquifer storage likely will be better for inter-annual storage needs. Managers can obtain efficient operating rules that are applicable to real control problems: given a method for determining the optimal control of a conjunctive-use system, control decisions can be reduced to simpler rules useful for practical operation of a system.

In practice, reservoirs are often operated according to rule curves that define desired storage and release targets [Young, 1967; Loucks *et al.*, 1981] that are often presented in graphical form as guides for those responsible for reservoir operation. Many rules are the result of negotiation between competing interests, and may be legally defined to protect operators and managers against liability or conflicts among beneficiaries and agencies. Johnson *et al.* [1991] noted that most models use heuristic guidelines to define the system's operating policy. Often these are developed for application to a specific system; however, a few general rules exist. For example, the "space rule" allocates reservoir storage space in proportion to expected inflows in a parallel water supply system. The similar "New York City rule" allocates reservoir storage space to obtain an equal spill probability for each reservoir [Johnson *et al.*, 1991].

Though optimization efforts can improve our management of conjunctive-use systems, they are still of limited practical benefit. Economic, social, and legal aspects of the conjunctive-use problem are still absent; as Burges and Maknoon [1975] state "There is no advantage in building or using a model of a conjunctive ground-surface water system that includes considerable hydrologic detail but neglects to adequately represent legal and economic nuances."

Also, Yeh [1985] notes that few practical rules have yet been developed through application of optimization techniques developed in the literature:

"reservoir operators are still very reluctant to use optimization models for their day to day scheduling of water releases and power generations." Often this is because (1) operators have not been directly involved in the development of the computer model, (2) most published papers deal with simplified reservoir systems and are difficult to adapt; and (3) there are institutional constraints that impede user-research interactions.

In addition, though conjunctive use potentially can improve the efficiency and reliability of a water supply system, stumbling blocks remain in its implementation. *Lettenmaier and Burges* [1979] point out that practical implementation of aquifer storage is inhibited by the manner in which costs of surface storage facilities are allocated, and that legal questions about the control of subsurface storage remain to be answered.

III. Modeling Issues

In any modeling effort, we generally must make simplifying assumptions about system characteristics. These are required to develop a problem that we can solve. For dynamic systems, the modeling of uncertainty and the stage-wise evolution of the system present particular challenges. In addition, we also must consider factors that affect the efficiency and tractability of the solution algorithm used to solve a conjunctive-use problem.

A. QUANTIFYING SYSTEM CHARACTERISTICS

The general application of optimization techniques to water management problems has resulted in some topics of controversy. Here I consider a few of these topics and how they should influence our optimization efforts.

In practice, there are a number of potential difficulties that we must consider when modeling water systems. Two difficulties have been especially controversial: (1) quantifying the costs associated with management decisions; and (2) characterizing the probability distributions of stochastic components in the system. In addition to these two difficulties, we also must select state variables cautiously to ensure a tractable problem that is still sufficiently realistic to be of practical benefit.

1. The Objective Function

Often managers must consider costs which are poorly defined. Estimated costs of rationing and floods are approximate at best, and may neglect important impacts. Estimated costs for other values may be entirely absent: costs attributed to environmental effects, other third-party effects (i.e., externalities), risk aversion, recreation, aesthetics, etc.

Nevertheless, though these factors are difficult to quantify, there generally are clear qualitative differences. Though we may not know the precise costs associated severe flooding, we know that costs increase with the severity.

To obtain quantitative measures of cost, we often use values provided by market transactions. For example, revenue obtained from the sale of hydropower may be an appropriate measure of benefit. However, though market values may be convenient and widely accepted, they may be inappropriate when markets do not incorporate important externalities.

No Specified Costs

Yao and Georgakakos, [1993] observe that "Reservoir systems are usually multiobjective and their operation involves a wide spectrum of considerations including interagency coordination, legal and regulatory issues, and public pressures, among others. Such criteria cannot be modeled and, therefore, an operation which maximizes or minimizes one particular objective is meaningless and is not practiced." Instead, they advocate a set approach that successively limits available decisions in order to maintain the system within constraints.

However, *Yao and Georgakakos* recognize that their approach is limited, primarily because it fails to distinguish the desirability of different feasible solutions. Also, for a long enough time horizon, they state that there may be no feasible solution. While their approach can define the effect that different constraints have on the remaining management options (and therefore can be useful to policy makers in determining the constraints to apply to operation of a system), it cannot tell a manager what the optimum control strategy should be. Nor can it tell the operator what real-time control decisions to implement. In addition, constraints used by the set approach to establish "feasible" controls are not truly bounds on feasible solutions; often they are merely a result of management decisions or of negotiation.

Hybrid Specification of Costs

In contrast to the work of *Yao and Georgakakos*, most efforts to optimize water system management assume that an objective function can be developed to determine optimum control strategies. Often, however, these efforts use a hybrid approach that avoids a complete definition of costs.

When confronted with poorly defined costs or other difficult to model criteria, it has been common for model developers to use the extreme assumption that nothing is known about these costs or effects. Often this approach is taken to avoid an explicit statement of costs, such as that associated with the loss of human life, because of the ethical and political problems associated with such statements.

However, management decisions always contain an implicit definition of costs, even for those criteria for which we have not directly specified costs. For example, chance constraints are often used when there is significant uncertainty in the costs or probabilities associated with extreme events. In flood control problems, the use of chance constraints allows us to avoid quantifying the costs

associated flood events that exceed some threshold of chance. This, however, ignores our knowledge that larger floods cause greater damage, in spite of exceeding a threshold of chance. The use of chance constraints results in an inferior objective value that considers the cost of all extreme events to be the same. Similarly, the solution boundaries established by the set approach of Yao and Georgakakos, [1993] are equivalent to specifying an objective which has a zero cost for "feasible" solutions and one which has very large costs for "infeasible" solutions.

Specified Costs

Any approach that ignores information or incorporates it in the objective by a prescribed approach will generally perform more poorly than one that makes appropriate use of this information, no matter how poorly quantified.

On the other hand, some methods have been developed to provide a measure of non-market costs and benefits. One of the more popular is the contingent valuation method (for discussion and a recent application, see Whittington *et al.*, [1993]). Though these other methods can contribute to the solution of the conjunctive-use problem presented in this thesis, I assume that all costs associated with system operations are quantifiable and that the cost function used to evaluate the management solution is known. The difficulties associated with including difficult to quantify (often non-market) costs is not considered here.

2. The Probability Distribution of Stochastic Inputs

Describing the continuous probability distributions of stochastic components, such as stream flow, is the second difficulty encountered in modeling efforts.

In many systems, the amount of information is insufficient to determine a unique model for a probability distribution. The best model will often depend on the problem under consideration. For certain types of problems, the reliability of a control solution may be poor. Even for the same data set, such as historical stream flows of a particular river, a control solution for extreme events such as floods will be less reliable than one for mean characteristics such as water supply. The probability of extreme events, such as of floods and droughts, is almost never known with much accuracy. Even our knowledge of the mean and variance of many stochastic elements can be very poor.

Because of this difficulty, some authors have advocated approaches to uncertainty that may use less knowledge of a distribution than we have available [Yao and Georgakakos, 1993]. The problem with these approaches is similar to that of quantifying an objective function; using less knowledge does not improve our solutions, it merely avoids decisions that may be awkward. Others have suggested methods to creatively incorporate available information without developing a continuous model of a distribution [Kelman *et al.*, 1990]. Even if we decide that it is best to use a probability model, selecting an appropriate model can be difficult [Vogel and Fennessey, 1993].

For problems in which managers are still unwilling or unable to quantify certain costs or distributions, various approaches exist. *Foufoula-Georgiou and Kitanidis* [1988] use chance constraints with their GDP method. *Kelman et al.* [1990] avoids using a model of flow distribution by a sampling stochastic dynamic programming method that preserves many of the important characteristics of a distribution. *Yao and Georgakakos*, [1993] use a set approach that avoids the use of both an explicit objective and a model of flow distribution, though it cannot be used for real-time control except perhaps during emergency operations.

3. Incorporating Uncertainty

There are many sources of uncertainty in system modeling. I discuss here two sources of uncertainty that can be modeled by using stochastic variables. Both affect the transition of the system from one period to the next. The first is uncertainty in system inputs, such as stream flow. The second is uncertainty in the values of state variables used to describe the system. Though both are a result of limited knowledge about the system, the first is generally attributed to the stochastic, uncontrollable nature of input process, and the second is generally attributed to uncertainty in measurements or decisions over which some control exists.

Three cases of uncertainty are discussed below followed by an illustration of how they might be included in a system model. The first two deal with stochastic uncontrollable inputs, one that results from a random uncorrelated time series and one from an autocorrelated series. The third deals with an uncertain state variable where uncertainty can be reduced by more precise measurements.

Based on the previously discussed mathematical form of the model, uncertain inputs and uncertain measures of a system's characteristics affect control decisions by making the end-of-stage value of state variables uncertain. Optimum decisions minimize the sum of current and future costs, and future costs depend on the values of y . Again, the mathematical form of the control problem is

$$F_{(t)}(x) = \min_u \{ C(x,u) + \sum_k \{ p_k F_{(t+1)}(y_{(k)}) \} \}$$

where $y_{(k)} = T(x,u,w_{(k)})$ is a function of the current state x , decisions u , and stochastic variables w .

In the first case, the state of the system is influenced by a random uncorrelated input, a stochastic variable w_j is added to the vector w to represent the realization of this input during the current stage. Though the value of w_j is uncertain, its probability distribution is assumed known and capable of being modeled. Since the input is uncorrelated with prior realizations or other characteristics of the system, no additional state variables need be added to x .

In the second case, the state of the system is influenced by a random input that is autocorrelated, one or more state variables must also be added. This is necessary because the realized values of past inputs influence optimal decisions. The added state variables represent prior realizations of the input; however, if the model of autocorrelation requires many prior stream-flow values, much of the information could be incorporated by a single state variable representing a single stream-flow forecast [Stedinger *et al.*, 1984]. These state variables condition the distribution of the random variable, and discrete realizations used to span this distribution are specified as functions of these state variables. For certain simple models of probability distribution and autocorrelation, the effect of the state variables on the stochastic variable could be built into the transfer equation and the stochastic variable used to model only the independent residual variation.

In the third case, the value of a state variable has a significant amount of uncertainty, and a stochastic variable can be included to incorporate this uncertainty. Though not a physical input to the system, the stochastic variable represents additional uncertainty in identifying the effect of control decisions on the state of a system. If this uncertainty is due to measurement error, we may have some control over this uncertainty by making additional measurements. In

a dynamic system, measurements made over time present an uncertain time series that may also be autocorrelated or may otherwise depend on state variables. For example, if measurement errors of a state variable change with the variable's magnitude, then the stochastic variable representing this error depends on the state variable.

In each of the above cases, adding stochastic variables allows the specification of a transfer function that has a precise mathematical form similar to that discussed in Chapter II. Assuming the system is linear, the transfer equation can be modeled by a system of equations

$$y = T(x,u,w) = Ax + Bu + Cw$$

[Yao and Georgakakos, 1993]. Suppose a water supply system is described by two state variables where x_1 = known reservoir level, x_2 = either the prior month's stream flow or a snowpack measurement (Kelman *et al.*, 1990). If x_2 is the prior month's stream flow, we can incorporate autocorrelation by a first-order autoregressive model. If x_2 is a measurement of snowpack, we can obtain stream-flow estimates based on the correlation between snowpack and snow-melt runoff. Either of these values is subject to measurement error, but snowpack measurements are much more imprecise and subject to revision by later measurements. Given that there is one control decision for this system, u_1 = water release to meet current demand, the resulting state y depends on u_1 and $w_1 = w_1(x_2)$ = current month's uncertain stream flow.

Consider first the model with x_2 = the prior month's stream flow. The transfer equation can be modeled as

$$y = T(x,u,w) = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} x + \begin{bmatrix} -1 \\ 0 \end{bmatrix} u + \begin{bmatrix} 1 \\ 1 \end{bmatrix} w$$

In this case, the uncertain input w_1 affects both end-of-stage state variables, y_1 and y_2 . Also, the dependence of w_1 on x_2 could be model as first-order autoregressive, where

$$w_1 = \mu + \rho(x_2 - \mu) + \varepsilon\sigma\sqrt{1-\rho^2}.$$

Because of the simple model of stream flow autocorrelation, the dependence of current stream flow on prior realizations could be included in the transfer function as

$$y = T(x, u, w) = \begin{bmatrix} 1 & 0 \\ 0 & \rho \end{bmatrix} x + \begin{bmatrix} -1 \\ 0 \end{bmatrix} u + \begin{bmatrix} 1 \\ \sqrt{1-\rho^2} \end{bmatrix} w$$

where ρ is the sample lag-one autocorrelation coefficient. Using this model, w_1 is now independent of x . A problem modeled with this transfer function could be solved by the algorithm of *Foufoula-Georgiou and Kitanidis [1988]* without extension.

Alternately, consider the model with x_2 = a measurement of snowpack. Vector w now contains two stochastic variables: w_1 = current month's uncertain stream flow, and w_2 = random error in the season's stream-flow estimate. For this system the transfer equation can be modeled as

$$y = T(x, u, w) = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} x + \begin{bmatrix} -1 \\ 0 \end{bmatrix} u + \begin{bmatrix} 1 & 0 \\ -1 & 1 \end{bmatrix} w .$$

In this case, the uncertain input w_1 still affects both end-of-stage state variables, y_1 and y_2 . Also, w_1 still depends on the state of the system as some function of x_2 . In addition, the random error w_2 now affects the value of state variable y_2 , but is independent of the values of the state variables x .

Both the prior realization of stream flow and the stream-flow estimate for the snow-melt season could be used in the prediction of current stream flow; however, this would be at the cost of an added dimension to the state vector. For example, let x_1 = known reservoir level, x_2 = the previous month's stream flow, and x_3 = a measurement of snowpack. Vector w is defined as in the previous paragraph. The transfer equation can be modeled as

$$y = T(x, u, w) = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix} x + \begin{bmatrix} -1 \\ 0 \\ 0 \end{bmatrix} u + \begin{bmatrix} 1 & 0 \\ 1 & 0 \\ -1 & 1 \end{bmatrix} w .$$

In this case, the distribution of w_1 now depends on a model that includes both x_2 and x_3 .

4. Modeling Stochastic Variables

The distribution of each stochastic variable is modeled by choosing discrete realizations $w_{(k)}$ of the multi-variate distribution w , and with each $w_{(k)}$ we can specify an associated probability. For example, if a stochastic variable w_j

is normally distributed, its distribution could be approximated by discrete values $w_{j,(k)}$ at

$$\sigma = 0, \pm 1, \pm 2, \dots$$

standard deviations from the mean. These realizations could then be used to approximate the distribution over the intervals

$$\dots, [-2.5, -1.5], [-1.5, -0.5], [-0.5, 0.5], [0.5, 1.5], [1.5, 2.5], \dots$$

with probabilities

$$p_\mu = 0.383, p_{\mu \pm 1\sigma} = 0.242, p_{\mu \pm 2\sigma} = 0.061, \dots$$

which are the weights applied to each realization when calculating expected values. Derivatives $\partial w_{(k)} / \partial x$ are estimated by changes in the k 'th ordered discrete realization $w_{(k)}$ while maintaining the same deviation from the mean. Given some simplifying assumptions discussed below, this matrix could also be approximated by the covariance matrix $\text{Cov}\{w, x\}$.

Determining Derivatives

It may be possible to calculate the elements of this matrix analytically using the probability distribution model for w . In the model developed in Chapter II, I was able to use the formula

$$\frac{\partial w}{\partial x} = \begin{bmatrix} 0, & 0, & \frac{w_1 - \chi_\tau}{x_3 - \chi_{\tau-1}} \rho_\tau \frac{\sigma_\tau}{\sigma_{\tau-1}} \end{bmatrix}$$

based on the three-parameter log-normal model developed.

In some situations analytic determination of derivatives may be difficult. As an alternative, each value $\partial w_{j,(k)} / \partial x_{j,(i)}$ can be approximated by regressing w_j on x_j from historical data. Assuming that changes in x_j have the effect of shifting the distribution of w_j without changing the shape (an appropriate simplifying assumption with the amount of data normally available), $\partial w_{j,(k)} / \partial x_{j,(i)}$ will be equal to the slope of the regression equation at $x_{j,(i)}$ regardless of realization $w_{j,(k)}$ (Figure III-1). If historical data are limited, only a simple linear regression of w_j on x_j may be possible. In this situation, $\partial w_{j,(k)} / \partial x_{j,(i)}$ will be constant for any $x_{j,(i)}$ and any $w_{j,(k)}$. Matrix $\partial w / \partial x$ is then

a constant matrix $\mathbf{R}$ which could be calculated just once at the beginning of the problem.

Such a matrix $\mathbf{R}$ could be calculated from covariances as follows [Draper and Smith, 1981, Chapter 2]

$$r_{jj} = \text{Cov}(w_j, x_j) / \sqrt{\text{Var}(w_j) \text{Var}(x_j)}$$

and is related to the matrix $\mathbf{P}$ of correlation coefficients between $\mathbf{w}$ and $\mathbf{x}$ where the terms of $\mathbf{P}$ are calculated as follows

$$\rho_{jj} = \text{Cov}(w_j, x_j) / \sqrt{\text{Var}(w_j) \text{Var}(x_j)} .$$

Thus, the matrix $\partial \mathbf{w} / \partial \mathbf{x}$ preserves the correlation between $\mathbf{w}$ and $\mathbf{x}$ and the effect of predictors on the stochastic variables.

Other alternatives to the above approach exist. If correlation between stochastic inputs and state variables can be explicitly modeled, an alternate approach may be to incorporate this knowledge directly in the transition function. This leaves only uncorrelated residual uncertainty to be modeled by stochastic variables $\mathbf{w}$, and $\partial \mathbf{w} / \partial \mathbf{x} = 0$.

5. Incorporating Trends, Discount Rates, and Seasonality

It is possible to also incorporate various other system characteristics into this approach, such as trends, discount rates, and seasonality. These characteristics could produce changes in the cost function, transition function, or constraints.

Trends can be incorporated by allowing the current cost function to change with time, though the change would have to be applied in the reverse of its actual development since the recursions go backwards in time. The recursive update of the expected cost function would then take the form

$$F_{(t)}(\mathbf{x}) = \min_{\mathbf{u}} \{ C_{(t)}(\mathbf{x}, \mathbf{u}) + \sum_k \{ p_k F_{(t+1)}(\mathbf{y}_{(k)}) \} \}$$

where $C_{(t)}$ varies with each recursion. Similarly a trend could be applied through changes in a constraint applied to each recursive solution. Such an approach would allow, for example, changing water availability resulting from use of senior reserved water rights, such as water reserved for area-of-origin or for federal lands.

A discount rate of α percent can be applied by multiplying the updated expected cost function by the value $\beta = 100 / (100 - \alpha)$ in each recursion to reduce the present value of future costs:

$$F_{(t)}(\mathbf{x}) = \min_{\mathbf{u}} \{ C(\mathbf{x}, \mathbf{u}) + \beta \sum_k \{ p_k F_{(t+1)}(\mathbf{y}_{(k)}) \} \} .$$

The effect of a discount rate applied to the entire planning period can be more easily observed in a deterministic problem where the cumulative effect of a discount rate over the entire planning horizon would be represented by

$$F_{(0)}(\mathbf{y}) = \min_{\mathbf{u}} \{ \sum_{t=1}^N \beta^t [C(\mathbf{x}_{(t)}, \mathbf{u}_{(t)})] \} = \sum_{t=1}^N \beta^t [C(\mathbf{x}_{(t)}, \mathbf{u}_{(t)}^*)] .$$

In a deterministic problem, $[\mathbf{u}_{(1)}^*, \dots, \mathbf{u}_{(N)}^*]$ is the single control schedule needed to specify the solution for all stages.

Seasonality can be incorporated by a variety of approaches, such as by developing a unique expected cost function for each season or by adding another variable to the state vector. While the use of a state variable is more consistent with the definition of the state of a system, it requires development of a variable whose value is periodic. Oscillating functions such as a sine or cosine could be used, but they may imply intra-seasonal relationships that do not exist. Seasonality is most often incorporated by developing a unique expected cost function for each season. The backwards stage-by-stage solution of an expected cost function develops as previously discussed, but costs, constraints, stochastic distributions, or even the transition function may vary periodically with the current stage's season. We can represent this by

$$F_{(t)}(\mathbf{x}) = \min_{\mathbf{u}} \{ C_{(\tau)}(\mathbf{x}, \mathbf{u}) + \sum_k \{ p_{k, \tau+1} F_{(t+1)}(\mathbf{y}_{(k)}) \} \}$$

subject to constraints of form $\mathbf{g}_{(\tau)}^T \mathbf{u} \leq d_{(\tau)}(\mathbf{x})$

where τ is an index for the season and is specified uniquely for any stage t . Note that $F_{(t)} = \min_{\mathbf{u}} \{ f_{(t+1)} \}$ regardless of the season. For sufficiently long planning horizons, control solutions still converge, but only when solutions for equivalent seasons are compared. For example, if a monthly decision period is used as in Chapter II, convergence of control decisions is observed when comparing the most recent twelve recursions with the previous twelve.

Regression of Stochastic Variables on State Variables




Fig. III-1. Regression of w_j on x_j' and underlying probability distribution.

B. COMPROMISES

In addition to difficulties in quantifying system characteristics, modeling also requires compromises between accuracy and feasibility. Here I present two characteristics of the conjunctive-use problem that required consideration in developing the problem presented in Chapter II.

1. Selecting State Variables

As *Stedinger et al.* [1984] point out, the better the hydrologic state variables selected to predict stochastic characteristics of the system, the better control decisions will be. However, as they also point out, if all information on the state of the system is included, the dimension of x can become large. It is necessary for us to choose predictors that transmit as much information about a stochastic process as possible with the fewest number of state variables possible.

This can often be accomplished by representing several characteristics of the system by a single variable. For example, *Stedinger et al.* [1984] suggest that an estimate of future stream flow can be useful in summarizing information about prior flows. *Kelman et al.* [1990] use an estimate of the snow-melt season's runoff as a state variable, switching to the prior month's flow when this provides better information (and even dropping the extra state variable entirely for those months when month-to-month flow correlations were modest).

Incorporating Autocorrelation

When prediction is based primarily on autocorrelation, traditional autoregressive and moving average models can be employed. A practical description of observed autocorrelation will rarely require the addition of more than one or two state variables. *Chatfield* [1989] provides a good practical guide in fitting a model, and *Box and Jenkins* [1976] provide an in depth analysis of the popular ARIMA model.

Another approach to including the information of predictors without using extra state variables is the adaptive-control method [*Bras et al.*, 1983; *Stedinger et al.*, 1984]. This approach does not include all predictors in the state vector but instead accounts for the effect of these predictors by making transient, non-stationary modifications of distributions of the stochastic variables. A major advantage of using adaptive control is that it reduces the number of state variables needed to describe the state of the system. A main disadvantage, however, is that it does not allow development of a global solution. Adaptive-

control develops only the solution which is appropriate for a single state of the system and for particular values of the predictors. Thus a new solution must be developed for each period and each change in the predictors.

2. Length of a Stage

Stages are increments in time or space identified as a unit for the application of control decisions. For example, a real-time management problem requires that control decisions be updated periodically; a process control problem requires that decisions be made for each section of the assembly line or treatment path. Generally, a stage is defined as the maximum time that a system can be operated without updating control decisions. Stages can also identify spatial lengths such as when optimizing the steady-state degradation of effluent down a stream channel [Cardwell and Ellis, 1993].

Each stage is represented by one recursion of the solution. The appropriate length of a stage varies with the problem being considered. Generally, short stages are desirable for more accuracy; however, there are often practical limits on how short a stage can be, and the length of a stage is a compromise between accuracy and practicality.

The Improved Efficiency of Short Stages

The length of a stage affects solution accuracy in several ways. Stochastic inputs to the system make the cost of decisions uncertain. Inaccuracies introduced by this uncertainty can frequently be reduced by choosing a short stage. As shorter stages are chosen, current costs $C(x,u)$ become negligibly different from the expected value $E_w\{C(x,u,w)\}$. This is when inputs are highly autocorrelated.

For example, in a water supply system with a random demand component, the actual loss incurred from a mismatch between demand and a release decision will not be precisely known until the end of the stage. However, because demand for water is highly autocorrelated—especially on short time scales—the shorter the stage used, the smaller the error in the cost estimate.

Another related consideration is the inefficiency that may result from applying constant decisions to a long stage. However, generally the state of a system changes gradually, and thus the optimum decisions also change slowly. The shorter the stage used, the smaller the mismatch between the fixed decisions that are applied and the optimum decisions.

Practical Limits on Short Stages

There are a variety of practical limits on using short stages. One universal difficulty is that shorter stages impose a greater computational burden in solving the problem. There may also be physical limits of the system that prevent rapid implementation of changes (such as adjustment of releases controlled by weirs) or that prevent frequent updates of state variables (such as efforts to obtain field measurements of snow-pack accumulation). Though advances in technology and system infrastructure could address many limitations (such as using remote sensing for snow-pack measurements), until such technology is in use, it is unproductive to use a stage that is shorter than these limitations allow.

The appropriate stage length will also depend on the characteristics of system being modeled. For a water supply system, random fluctuations in demand over a period of months may be significant but over a period of one week may be insignificant. Short term fluctuations become even less significant if there are significant storage components which are not identified by a system model. For example, a water distribution system has storage capacity downstream of reservoirs; though this storage may not be included in a system model, it still can contribute to optimal control by reducing short-term demand fluctuations.

A special difficulty can exist when solving multi-objective problems with mixed time scales. For example, the stage length required for considering flood control objectives may be inconveniently small for multi-objective problems that also consider water supply or hydropower generation. In addition to the difficulty of finding an appropriate stage length, it may also be difficult to define a model of stream flow that preserves the characteristics needed by different objectives.

IV. Solution of Lumped Parameter Models Containing Autocorrelation by Gradient Dynamic Programming

A. INTRODUCTION

Real-time control problems are a type of dynamic problem arising in many operational environments such as process control and resource allocation. Often these problems are characterized by uncertain inputs. For such problems, discrete dynamic programming (DDP) is the only general method that can determine optimal control solutions without major and often inappropriate simplifying assumptions. Unfortunately, DDP is limited in its ability to solve complex problems due to the exponential growth of computational effort, referred to as the "curse of dimensionality". As a result, the discrete dynamic programming (DDP) approach has been limited to solving lumped parameter problems consisting of only two or three state variables.

Foufoula-Georgiou and Kitanidis [1988] developed the new method of Gradient Dynamic Programming (GDP) to reduce this limitation and solve problems consisting of perhaps six or more state variables. Though we can solve more complex problems using GDP, we cannot include autocorrelation of inputs, such as stream flow, in the existing algorithm. I present an extension of the GDP method that allows its application to control problems that include autocorrelated inputs or other information that can be used to improve estimates of uncertain inputs.

This section presents justification for the use of dynamic programming and for the development of GDP as an improved method for solving dynamic programming problems. Section 2 outlines the basic GDP approach and Section 3 presents the extended GDP algorithm. The final section discusses general considerations in selecting GDP as an optimization method.

1. Discrete Dynamic Programming

Dynamic programming describes a class of optimization techniques applied to the control and planning of dynamic systems [*Yakowitz, 1982*]. Dynamic systems are those that progressively change through stages, representing increments in either time or space. Though frequently confused with the more general concept of dynamic programming, DDP is but one class of methods available for solving dynamic programming problems. DDP has long

been a popular optimization technique, especially in water-systems management [Yakowitz, 1982; Yeh, 1985]. The popularity and success of this technique is primarily due to its ability to incorporate nonlinear constraints and objectives, and its ability to also incorporate stochastic variables [Willis and Yeh, 1987]. Indeed, DDP remains the only approach of any universality in solving dynamic problems containing stochastic inputs [Yakowitz, 1982]. This is in contrast to deterministic dynamic programming problems, which can be solved by a variety of approaches.

Because dynamic problems containing stochastic inputs present special challenges, they are often identified as stochastic dynamic programming (SDP) problems. Most hydrologic systems, for example, present SDP problems because they contain stochastic inputs such as stream flow and rainfall. Since DDP is the only approach generally applicable to the solution of SDP problems, SDP is commonly synonymous with DDP. However, unlike DDP, SDP does not specify a particular optimization approach. A variety of approaches have been used to solve specific SDP problems [Karamouz and Vasiliadis, 1992], though many, including GDP, are variants of DDP.

DDP solves a control problem by determining decisions that minimize expected cost or other objective given knowledge of the current "state" of a system. The "state" of a system is a general term that describes any changing characteristics that influence optimal decisions. Both the state of the system and optimal decisions evolve from one stage to the next as a result of prior decisions and uncertain inputs.

Unlike deterministic systems, a single control schedule does not define future states of a stochastic system. Because future control decisions depend on the uncertain evolution of future states of the system, a stochastic solution requires the generation of control decisions and expected costs for every possible state of the system and each stage [Kelman *et al.*, 1990].

Because most systems are characterized by continuous state variables, there are an infinite number of unique states; and the solution consists of an infinite number of control decisions. One way that we can describe an infinite number of control decisions is to assume that the expected cost and control decisions have a particular form of functional dependence on the state variables. For example, Karamouz and Houck [1982] specified decisions as a linear function of state variables, and Buras [1972] recommended using a fitted polynomial function of high order. This approach generally results in sub-optimal solutions,

though it is frequently used to provide a local approximation in deterministic solution methods, such as Differential Dynamic Programming [Yakowitz, 1982; Jones et al., 1987; Trezos and Yeh, 1987, Culver and Shoemaker, 1993].

In contrast, DDP methods summarize an operating policy by identifying the expected costs and optimal control decisions at each of a finite number of discrete states and interpolating values at intermediate states. To optimally control a dynamic system using DDP, an objective function value is calculated for each feasible transition between states of each stage of a problem. The optimum control decisions are then chosen as those that produce the minimum expected cost.

A caution in using DDP is that one should be wary of proposed solutions until the deleterious effects of discretization have been somehow bounded and found acceptable [Yakowitz, 1982]. As Gal [1979] pointed out, a DDP approximation is optimal only for the discretized version of the model. In spite of the simplifying assumptions required to use a continuous, global function such as proposed by Buras, a DDP approximation can perform worse if the appropriateness of the discretization is not verified.

Because dynamic programming allows us to avoid specifying in advance the functional relationship between the solution (optimum decisions) and the independent variables (state of the system), almost any form of non-linearity in this relationship is possible. This ability is especially useful in studying problems where prior work has not provided a basis for assuming simple pre-determined functional relationships.

2. Gradient Dynamic Programming

Unfortunately, DDP problems become computationally intractable if many state variables are required to describe the state of a system. DDP describes each continuous state variable by a number of discrete values, and each combination of these values is assessed in finding a solution. As a result, the effort required to solve a problem is $T(n) V^n$, where n is the number of continuous state variables (referred to as the "dimension" of the problem), V is some average number of discrete values used to span each variable, and $T(n)$ is the effort to determine a solution for each of the V^n discrete states. With each additional variable added to a DDP problem, the number of unique combinations is multiplied by the number of discrete values used to span the new variable. Esmail-Beik and Yu [1984] provide a good illustration of this

difficulty in their solution of a reservoir management problem. This "curse of dimensionality" inherent in the solution of DDP problems has generally limited its practical application to problems described by only two or three state variables [Yakowitz, 1982; Johnson *et al.*, 1993].

There are two basic approaches available to reduce the computational effort required to solve a DDP problem. The first is to reduce the number of discrete values used to span the range of the state variables. This approach may lead to a deterioration in solution accuracy, however. The second is to reduce the number of variables needed to describe a system [Gal, 1979; Saad and Turgeon, 1988; Saad *et al.*, 1992]. For example, Kelman *et al.* [1990] modeled the storage of a multi-reservoir system with a single state variable for storage. However, aggregation of characteristics can be inappropriate if we are to represent the system realistically. Beyond these two general approaches, Johnson *et al.* [1988, 1993] discuss a variety of other techniques that may significantly reduce the computational effort of some problems. These include solving a sequence of problems with increasing difficulty, eliminating from consideration uninteresting areas of the state-space, partitioning the original problem into smaller separable problems, or creatively choosing variables used to model the system.

To allow application of dynamic programming techniques to problems of larger dimensionality, Kitanidis and Fouloula-Georgiou [1987] and Fouloula-Georgiou and Kitanidis [1988] took the first basic approach in developing GDP. GDP is able to significantly reduce the number of discrete values required to span state variables while still achieving a reasonable level of solution accuracy. They accomplished this by using a more accurate interpolation method to determine the values cost and decision values between discrete states. Though the effort required to solve a problem using GDP still increases at an exponential rate V^n with dimension n of the problem, the base of the exponent V can be reduced. Thus one can greatly reduce the number of unique combinations of discrete values that must be considered.

GDP is able to employ a more accurate interpolation scheme by using an interpolation that is smoother and that preserves additional information at each discrete state. In addition to preserving the values at each node, GDP also determines and preserves the gradients of these values with respect to the state variables. These gradients can be calculated analytically (as presented here) or approximated by a finite difference approach. To interpolate a function between discrete states while preserving both the values and gradients, Kitanidis and

Foufoula-Georgiou [1987] proposed using Hermite interpolation, developed for this purpose by *Kitanidis* [1986]. They demonstrated that with decreases in the discretization interval Δx , the error of the control policy and the cost functions converge as $(\Delta x)^3$ and $(\Delta x)^4$, respectively, using Hermite interpolation versus Δx and $(\Delta x)^2$ using linear interpolation.

3. Extension to Include Autocorrelation

A main purpose of this work is to extend the GDP algorithm to allow its application to problems with autocorrelation.

To model autocorrelation, one frequently includes prior realizations of stochastic inputs as state variables. The probability distribution of stochastic inputs is then an unchanging or "stationary" function of the state of the system.

Many previously solved stochastic control problems have ignored autocorrelation because of the limited number of state variables that DDP methods have been able to handle. Control problems that have included autocorrelation have been greatly simplified. Most frequently these simplifications have involved the reduction of significant portions of a system to description by a single state variable [*Gal*, 1979; *Kelman et al.*, 1990]. Various authors have also applied adaptive-control methods to incorporate autocorrelation while avoiding the use of extra state variables [*Bras et al.*, 1983; *Stedinger et al.*, 1984; *Karamouz and Vasiliadis*, 1992]. Adaptive-control incorporates the effect of autocorrelation by making transient modifications to the probability distributions of the stochastic inputs. As a result, adaptive-control does not develop global solutions: it only solves the optimal control for a single state of the system and for a particular set of prior inputs. Thus a new solution must be developed for each different state, prior input, or stage.

Compared to other DDP methods, GDP has greater flexibility to include extra state variables; however, the existing GDP algorithm is not general enough to describe dependence of stochastic inputs on state variables and thus is unable to incorporate autocorrelation. The algorithm presented here extends the GDP to incorporate this dependence. In addition to autocorrelation, the extended GDP algorithm can also incorporate other correlated information that conditions the probability distribution of stochastic variables.

B. PROBLEM SPECIFICATION

Dynamic control of a system requires that managers determine control decisions for each of a series of stages that define a planning horizon or a process. For example, real-time management requires that managers determine periodic decisions about how to operate the system.

When decisions are influenced by uncertain inputs, a manager cannot specify effective and efficient decisions by a single control schedule. Prior values of stochastic inputs influence optimal decisions, and--without foreknowledge--a manager does not have all information needed to determine optimal decisions for a stage until the beginning of that stage.

Information that influences optimal decisions and that describes changing system characteristics specifies the state of a system. A solution of an optimal control problem is the set of optimal decisions and expected costs for each stage and state of a system. We can describe these decisions and costs as a function of the state of the system. For example, optimal water releases from a reservoir and expected risk of water rationing or flooding can be described as a function of current reservoir levels, stream flows, and season.

1. Modeling a Control Problem

We can measure the performance of decisions with an objective function that, in its most general form, can be represented as any function

$$f = f(\mathbf{x}, \mathbf{u}, \mathbf{w}, \mathbf{y})$$

where $\mathbf{x}$ is a vector of independent state variables describing the system at the beginning of the current stage, $\mathbf{u}$ is a vector of decision variables applied during the stage, $\mathbf{w}$ is a vector of random variables whose values are realized during the stage, and $\mathbf{y}$ is a vector of state variables describing the system at the end of the stage. Vectors $\mathbf{x}$ and $\mathbf{y}$ describe the same state variables in adjacent stages; i.e., $\mathbf{y}_t = \mathbf{x}_{t+1}$ in adjacent stages t and $t+1$.

Variables, $\mathbf{u}$, $\mathbf{w}$, and $\mathbf{y}$ are implicit functions of state variables, $\mathbf{x}$. Autocorrelation is modeled by making the distribution of random variables dependent upon the state of the system, thus $\mathbf{w}(\mathbf{x})$. A transition function, $\mathbf{y} = \mathbf{T}(\mathbf{x}, \mathbf{u}, \mathbf{w})$, models the evolution of a system from state $\mathbf{x}$ to state $\mathbf{y}$. The control solution, $\mathbf{u}^*$, is specified as a function of $\mathbf{x}$, thus $\mathbf{u}^*(\mathbf{x})$.

In order to incorporate stochastic variables in a model, only a limited number of parameters can be used to model a probability distribution. We can parameterize a distribution either by using its moments or by using discrete probability-weighted realizations that span the distribution. In this work, I estimate w by discrete realizations $w_{(k)}$ with probability p_k . Likewise, because y depends on w , its values are uncertain and estimated by

$$y_{(k)} = T(x, u, w_{(k)}) \quad .$$

using the same probability weights.

Foufoula-Georgiou and Kitanidis [1988] developed the GDP method for stage-wise decomposable problems with the more limited objective function

$$f = C(x, u) + F(y)$$

that expresses the objective as a total "cost" composed of a current deterministic cost C and an expected future cost F . F is only a function of y since the state of a system specifies all information needed to determine expected costs of future decisions. For stochastic problems, y and $F(y)$ are uncertain and the objective is

$$f = C(x, u) + \sum_k \{p_k F(y_{(k)})\} \quad .$$

2. Solving a Control Problem

We can determine an optimum control solution by recursively solving the system

$$F_{(t)}(y) = f_{(t+1)}^*(x)$$

where

$$f_{(t+1)}^*(x) = \min_u \{ C_{(t+1)}(x, u) + F_{(t+1)}(y) \}.$$

is from the solution of the control problem in the previous recursion. The previous recursion solves the optimal decisions for the following stage, and the solution of the problem progresses backwards through time or space. The last stage is solved by the first recursion using an assumed function $F_{(N)}(y)$ representing costs expected beyond the end of the planning horizon.

If the planning horizon is very long or if future costs are discounted, a specified cost function $F_{(N)}(y)$ may have a negligible impact on current decisions and could be zero or any arbitrary value. On the other hand, if the cost function $F_{(N)}(y)$ significantly affects current decisions, it must have reasonable values. In practice, $F_{(N)}(y)$ commonly specifies expected costs that are the same for all final states; in other words, the state of the system at the end of the planning horizon doesn't matter. Less commonly, $F_{(N)}(y)$ specifies costs that vary with the final state. For example, if planned maintenance requires a certain final state, $F_{(N)}(y)$ could specify a high cost for any final state other than that required. In their study of a hydropower system, *Kelman et al.* [1990] used an assumed cost function based on the potential energy of water remaining in reservoirs.

In some well studied problems, we may know the general form of the function $F(y)$. For example, if we can assume that expected costs are a linear function of state variables, then $F(y)$ is defined by a few parameters that define the linear relationships. In many problems, however, we do not know a general form of $F(y)$. In these problems, we can use a DDP method that approximates $F(y)$ by discrete values spanning the state space. For each discrete state $x_{(i)}$, DDP solves

$$f^*(x_{(i)}) = \min_u \{ C(x_{(i)}, u) + \sum_k \{ p_k F(y_{(k,i)}) \} \}$$

where $y_{(k,i)} = T(x_{(i)}, u, w_{(k,i)})$, and $w_{(k,i)} = w_{(k)}(x_{(i)})$. Once we have determined optimal decisions $u^*(x_{(i)})$ and values of the objective $f^*(x_{(i)})$ for each $x_{(i)}$, interpolation between these values create piece-wise continuous functions for $u^*(x)$ and $f^*(x)$.

C. THE GDP METHOD

1. Notation

Vectors are column vectors. The transpose of a column vector is a row vector. Derivatives of a vector with respect to a scalar is a vector with the same dimension. The derivative of a scalar function with respect to a column vector is a row vector. Thus,

$$\mathbf{u}(x) = \begin{bmatrix} u_1(x) \\ u_2(x) \\ \vdots \end{bmatrix} \Rightarrow \frac{\partial \mathbf{u}}{\partial x} = \begin{bmatrix} \partial u_1 / \partial x \\ \partial u_2 / \partial x \\ \vdots \end{bmatrix}$$

but $f(\mathbf{x}) \Rightarrow \frac{\partial f}{\partial \mathbf{x}} = f_{\mathbf{x}} = [f_{x_1}, f_{x_2}, \dots]$

where $\mathbf{u}$ is a vector function of scalar x and f is a scalar function of vector $\mathbf{x}$.

It follows from these conventions that the derivatives of a m -dimensional vector function with respect to a n -dimensional vector is a $(m \times n)$ matrix of derivatives where the dependent vector function is indexed along rows and the independent vector indexed along columns. Second derivatives of scalar functions with respect to vectors are also matrices with the first vector indexed along rows and the second vector indexed along columns. Thus

$$\frac{\partial \mathbf{u}}{\partial \mathbf{x}} = \begin{bmatrix} \partial u_1 / \partial x_1 & \partial u_1 / \partial x_2 & \cdots \\ \partial u_2 / \partial x_1 & \partial u_2 / \partial x_2 & \cdots \\ \vdots & \vdots & \ddots \end{bmatrix} \text{ and } \frac{\partial^2 f}{\partial \mathbf{u} \partial \mathbf{x}} = f_{\mathbf{u}\mathbf{x}} = \begin{bmatrix} f_{u_1 x_1} & f_{u_1 x_2} & \cdots \\ f_{u_2 x_1} & f_{u_2 x_2} & \cdots \\ \vdots & \vdots & \ddots \end{bmatrix}$$

are matrices of derivatives, respectively, of vector function $\mathbf{u}(\mathbf{x})$ with respect to vector $\mathbf{x}$ and of scalar function f with respect to vectors $\mathbf{x}$ and $\mathbf{u}$.

When $f = f(\mathbf{x}, \mathbf{u}(\mathbf{x}))$, the definition of the total derivative with respect to $\mathbf{x}$ is

$$\nabla_{\mathbf{x}} f = \frac{df}{d\mathbf{x}} = \frac{\partial f}{\partial \mathbf{x}} + \frac{\partial f}{\partial \mathbf{u}} \frac{\partial \mathbf{u}}{\partial \mathbf{x}} = f_{\mathbf{x}} + f_{\mathbf{u}} \frac{\partial \mathbf{u}}{\partial \mathbf{x}}.$$

I distinguish total and partial derivatives of scalar functions, respectively, by the use of the subscripted symbol ∇ versus subscripts directly on the function. In contrast, total and partial derivatives of vector functions are distinguished by the use of 'd' and '∂' respectively, and will not use the symbol ∇ or subscripts.

2. Method

The following development of the GDP method parallels that of *Foufoula-Georgiou and Kitanidis* [1988] with the exception that I use index k to indicate a discrete realization of stochastic variables $w_{(k)}$ and index t to indicate a stage.

The goal is to provide rules that allow the system operator to make real-time operating decisions. These rules determine optimal decisions u^* that minimize an objective function f ; in other words

$$\min_u \{ f(x, u, w, y) \}$$

subject to the set of active constraints

$$Gu = d(x)$$

where G is a $(p \times m)$ scalar matrix of rank p , and $d(x)$ is a vector of the m right hand sides. The general form of an inequality constraint is

$$g^T u \leq d(x)$$

where g is a vector of coefficients and d can be a function of x . When binding, an inequality constraint becomes an equality constraint in the active set.

Using the method of Lagrange multipliers, optimal decisions that respect the constraints are also those that minimize

$$h(u, \lambda) = f + \lambda^T (Gu - d)$$

with respect to u and λ , where λ is a vector of Lagrange multipliers. Taking derivatives of the above equation, optimal controls are defined by the simultaneous equations

$$(\nabla_u h)^T = (\nabla_u f)^T + \lambda^T G = 0$$

$$(\nabla_\lambda h)^T = Gu - d = 0$$

Though these equations can verify optimal decisions u^* , they cannot be used to directly solve u^* . Objective function f is derived from the piece-wise continuous function F , and we do not know which "piece" of this function to use until after specifying the final state y resulting in part from control decisions.

Since evaluation of the above equations requires this knowledge, we cannot, in turn, determine optimal decisions until the appropriate "piece" of the expected cost function is specified.

Instead, given an initial set of control decisions, the algorithm determines incremental changes $\Delta \mathbf{u}$ that iteratively approach $\mathbf{u}^*$ by applying a Newton-type method. Subsequently, by determining $\mathbf{u}^*$ for each state $\mathbf{x}_{(i)}$, we can develop the implicit control function $\mathbf{u}^*(\mathbf{x})$ and cost function $f^*(\mathbf{x})$. Using $f^*(\mathbf{x})$ as the prior stage's expected cost function, $F(\mathbf{y})$, then allows us to determine a solution for the prior stage. This new solution then represents a planning horizon one stage longer.

3. Algorithm

The algorithm consists of three nested cycles, the most basic of which is the determination of $\mathbf{u}^*$. The optimal decisions $\mathbf{u}^*$ apply only to a discrete initial state $\mathbf{x}_{(i)}$ during recursion i . For each additional recursion, the algorithm solves $\mathbf{u}^*$ for every $\mathbf{x}_{(i)}$ until it has performed enough recursions to span a planning horizon of N stages or to allow control decisions to converge. Figure IV-1 illustrates the three cycles and the steps of the algorithm.

Step 1: Discretize the State Space

A control solution specifies $\mathbf{u}^*(\mathbf{x})$ and $f^*(\mathbf{x})$ as functions of state variables $\mathbf{x} = (x_1, \dots, x_n)^T$. Since we generally do not know the form of these functions, the algorithm solves piece-wise continuous functions by calculating values at a number of discrete states $\mathbf{x}_{(i)}$ and approximating intermediate values by interpolation.

The accuracy and efficiency of the GDP solution depends on the selection of an appropriate number of discrete values $\mathbf{x}_{(i)}$. Selection of discrete states involves trade-offs between solution accuracy and computational effort [Esmail-Beik and Yu, 1984]. Efficient placement of the $\mathbf{x}_{(i)}$ also requires judgment: state variables that induce rapid changes in a function need finer discretization; state variables that have a consistent effect on a function allow coarser discretization. We should verify a discretization's appropriateness prior to accepting a solution [Yakowitz, 1982] such as by observing the effect of using different discretizations.

Step 2: Specify an Expected Cost Function

The expected cost function $F(\mathbf{y})$ is a piece-wise interpolated function with known value F and gradients $\nabla_{\mathbf{y}} F$ at each discrete state $\mathbf{y}_{(i)}$. For the first

recursion, we must provide values and gradients that describe $F_{(N)}(y)$, the function for costs expected beyond the end of a planning horizon of N stages. For remaining recursions $t = N-1, N-2, \dots, 1$, the algorithm provides values and gradients from the previously solved recursion where $y_{(i),t} = x_{(i),t+1}$ and $F(y_{(i),t}) = f^*(x_{(i),t+1})$.

Given a discretization and values at each discrete state, GDP is able to interpolate values at intermediate points using Hermite interpolation (Appendix B). Hermite interpolation preserves both values and gradients at the discrete states that bound intervals [Kitanidis, 1986]. Fouloula-Georgiou and Kitanidis [1988, Appendix A] outline the Hermite interpolation method in greater detail, with correction that their equation (B7) should be $T = (1 - \eta_j)(1 - 3\eta_j)$ for $s = j$.

Step 3: Select a Discrete State and Initial Control Decisions

In order to develop functions $u^*(x)$ and $f^*(x)$, the algorithm determines optimal values for u^* , f^* , and $\nabla_x f^*$ for each discrete state $x_{(i)}$ by steps 4 to 9. Given an average number V of discrete values spanning each of the n state variables, the total number of discrete states is approximately V^n . We may be able to reduce this number and achieve greater solution efficiency by eliminating from consideration infeasible combinations of state variable values [Johnson et al., 1988, 1993].

With the specification of each discrete state $x_{(i)}$, we can select an initial control decision. By selecting initial decisions that are close to the optimum, we can improve the search efficiency and likelihood of convergence. Also, if the nature of the problem suggests initial decisions that satisfy all constraints, these decisions will be feasible and step 4 can be skipped. If no feasible decisions are apparent, step 4 will test a set of arbitrary decisions for feasibility prior to independently determining initial feasible decisions.

Step 4: Determine Initial Feasible Decisions

If the initial decisions specified in step 3 violate any of the constraints, the "phase 1" procedure of linear programming can determine feasible control decisions. Luenberger [1984] describes the phase-1 procedure.

Step 5: Calculate Improvement in the Decisions

The iterative improvement Δu is the solution to the system of equations

$$\begin{bmatrix} \nabla_u(\nabla_w f)^T & G^T \\ G & 0 \end{bmatrix} \begin{bmatrix} \Delta u \\ \lambda \end{bmatrix} = \begin{bmatrix} -\nabla_w f \\ 0 \end{bmatrix}$$

where

$$\nabla_w f = C_u + \frac{\partial y^T}{\partial u} \sum_k \{ p_k F_{y,(k)} \}$$

$$\nabla_u(\nabla_w f)^T = C_{uu} + \frac{\partial y^T}{\partial u} \sum_k \{ p_k F_{yy,(k)} \} \frac{\partial y}{\partial u}.$$

Matrix G is from the initial active set of step 4 or the updated set of step 6 or 7. Vector λ contains the p Lagrange multipliers needed to check the active constraint set in step 7. *Foufoula-Georgiou and Kitanidis* [1988] adapted this approach from the active set method outlined by *Luenberger* [1984, chapter 11]. The new solution is $u' = u_0 + \Delta u$ where u_0 is the vector of old decisions.

Step 6: Check Feasibility of Decisions

If an improvement Δu makes new decisions infeasible, the active constraint set is augmented with the violated constraint. The algorithm then identifies new feasible decisions u' by projecting u_0 on the domain defined by the updated active constraint set. The new control decisions are $u' = u_0 + \delta u$ where

$$\delta u = G^T(GG^T)^{-1}(d - Gu_0)$$

[*Foufoula-Georgiou and Kitanidis*, 1988, Appendix C].

An empty feasible decision space may result if the algorithm adds more than one constraint to the active set at a time. When an improvement Δu violates more than one constraint, the algorithm adds to the active set only the first constraint encountered when moving from u_0 in the direction Δu . The distance from u_0 to each new constraint is

$$\gamma = (d - g^T u) / g^T \Delta u$$

and the first violated constraint is that with the smallest distance γ .

Step 7: Check if Optimum Decisions are Found

Control decisions are close to optimum when the working set of constraints is still correct for the updated decisions and changes are small ($|\Delta u| < \delta$ or $\Delta u^T \Delta u < \delta$). If either condition is not met, the algorithm returns to

step 5 after removing any unnecessary constraints. Unnecessary constraints are those with a negative Lagrange multiplier ($\lambda < 0$).

Step 8: Calculate Values of Total Cost Function and Derivatives

Given the optimal decisions $\mathbf{u}^*$, the algorithm can now calculate the optimum value $f^*(\mathbf{x}) = \min_{\mathbf{u}} \{f(\mathbf{x}, \mathbf{u}, \mathbf{w}, \mathbf{y})\}$ and gradients $\nabla_{\mathbf{x}} f^*$ for the current initial state $\mathbf{x}_{(i)}$. These values are constants for the current $\mathbf{x}_{(i)}$ since the optimal decisions are also constants for $\mathbf{x}_{(i)}$.

The optimal expected value of the objective function is

$$f^*(\mathbf{x}_{(i)}) = C(\mathbf{x}_{(i)}, \mathbf{u}^*) + \sum_k \{ p_k F(\mathbf{y}_{(k,i)}^*) \}$$

where $\mathbf{y}_{(k,i)}^* = \mathbf{T}(\mathbf{x}_{(i)}, \mathbf{u}^*, \mathbf{w}_{(k,i)})$.

Calculation of the matrix of derivatives $\nabla_{\mathbf{x}} f^*$ requires $d\mathbf{u}^*/d\mathbf{x}$. The algorithm calculates the matrix of values $d\mathbf{u}^*/d\mathbf{x}$ by solving the system of equations:

$$\begin{bmatrix} \nabla_{\mathbf{u}}(\nabla_{\mathbf{w}} f)^T & \mathbf{G}^T \\ \mathbf{G} & 0 \end{bmatrix} \begin{bmatrix} d\mathbf{u}^*/d\mathbf{x} \\ d\lambda/d\mathbf{x} \end{bmatrix} = \begin{bmatrix} -\mathbf{D} \\ \nabla_{\mathbf{x}} \mathbf{d} \end{bmatrix}$$

where
$$\mathbf{D} = C_{\mathbf{ux}} + \frac{\partial \mathbf{y}^T}{\partial \mathbf{u}} \sum_k \{ p_k F_{\mathbf{yy},(k)} \left(\frac{\partial \mathbf{y}}{\partial \mathbf{x}} + \frac{\partial \mathbf{y}_{(k)}}{\partial \mathbf{w}_{(k)}} \frac{\partial \mathbf{w}_{(k)}}{\partial \mathbf{x}} \right) \}$$

and
$$\nabla_{\mathbf{x}} \mathbf{d} = \mathbf{d}_{\mathbf{x}} .$$

Appendix C provides the derivation of the above system of equations. Given $d\mathbf{u}^*/d\mathbf{x}$,

$$(\nabla_{\mathbf{x}} f^*)^T = (\nabla_{\mathbf{w}} f)^T \frac{\partial \mathbf{u}^*}{\partial \mathbf{x}} + C_{\mathbf{x}}^T + \sum_k \{ p_k F_{\mathbf{y},(k)}^T \left(\frac{\partial \mathbf{y}}{\partial \mathbf{x}} + \frac{\partial \mathbf{y}_{(k)}}{\partial \mathbf{w}_{(k)}} \frac{\partial \mathbf{w}_{(k)}}{\partial \mathbf{x}} \right) \} .$$

Step 9: Completing the Solution and Results

Given the optimal controls $\mathbf{u}^*$, the algorithm returns to step 3 to repeat steps 3 to 8 for each initial state $\mathbf{x}_{(i)}$. Once the optimal values of $\mathbf{u}^*$, $d\mathbf{u}^*/d\mathbf{x}$, f^* , and $\nabla_{\mathbf{x}} f^*$ have been found for each discrete $\mathbf{x}_{(i)}$, the algorithm can specify $\mathbf{u}^*(\mathbf{x})$ and $f^*(\mathbf{x})$ for any $\mathbf{x}$ using Hermite interpolation.

The piece-wise continuous functions $\mathbf{u}^*(\mathbf{x})$ and $f^*(\mathbf{x})$ are the solution to a control problem having a planning horizon that is now one stage longer than

that of the previous recursion. Given $u^*(x)$ and $J^*(x)$, the algorithm returns to step 2 to repeat steps 2 to 8 for each remaining recursion $t = N-1, N-2, \dots, 1$.

Unlike many optimal control algorithms that require a forward recursion to solve for the optimal controls from the expected cost function, the backwards recursion generates both $u^*(x)$ and $J^*(x)$ for the entire planning horizon; therefore, the solution is complete following the backwards recursion.




Fig. IV-1. Flow chart for the GDP algorithm

D. COMMENTS AND CONCLUSIONS

To solve a control problem using simulation or optimization requires that we develop a simplified model of a system. GDP makes it possible to solve DDP problems using models that are less simplistic than past efforts. Though this makes GDP an appropriate optimization method for many problems, it is not necessarily the best DDP technique for all problems. GDP is complicated to implement, and becomes useful only when it is not possible to solve a problem by other techniques without excessive simplification. Given the limitations of most DDP methods however, this occurs frequently.

There are some specific characteristics of GDP that we should consider prior to applying it to a problem. In addition to its added complexity, GDP should also be evaluated in terms of its efficiency and ability to converge on a solution. I discuss these briefly prior to presenting my conclusions.

1. Efficiency

Johnson et al. [1988, 1993] compared the efficiency of various spline interpolations against the Hermite interpolation used in GDP and against linear interpolation. Spline techniques of *Johnson et al.*, like the Hermite interpolation, seek to improve accuracy and smoothness of estimates; however, they differ in that they do not use additional information from gradients.

Both spline and Hermite interpolations demonstrated the ability to solve a 4-dimensional problem with less than 1% the effort and equivalent accuracy of a linear method. They also demonstrated that Hermite interpolation is more accurate than spline interpolation for the same discretization of the state space. On the other hand, Hermite interpolation requires extra computational effort to obtain gradients; and they determined that, for their problem, GDP was marginally less efficient than their spline interpolation method.

While both the Hermite and spline interpolations can significantly improve the accuracy and efficiency of a DDP approach, it is not clear which of these is better for a particular application. Though it is not a purpose of this work to compare the performance of these different approaches, a few observations are in order.

In general, the smoother the function and the greater number of discrete values used to span the state-space, the less valuable becomes information provided by gradients. I expect that Hermite interpolation is especially

appropriate for irregular functions (though stochastic problems are generally less irregular due to averaging process involved in determining expectations). I also expect that Hermite interpolation is more appropriate for functions that must be spanned by a few discrete states, as may be required when the number of state variables is large or when computational resources are limited. In an extreme case using just two discrete values to span the range of a state variable, Hermite interpolation should generally provide a better approximation of the function.

Though the accuracy of an interpolated function is improved more by the addition of a discrete state than by knowledge of a gradient, more gradients are generally obtained for an equivalent amount of computational effort. Whether calculating gradients or calculating adding additional discrete states is more efficient may be difficult to know in advance.

2. Convergence

Past efforts to approximate an expected cost function by analytic functions have often relied on polynomials of high order. However, fitting high-degree polynomials to data over the entire state space can induce severe oscillations in the fitted function. Optimization applied to such a fitted function may incorrectly seek out local optima induced by these oscillations [Johnson *et al.*, 1993]. Unlike a polynomial fit, DDP avoids this problem since it does not try to fit any predetermined functional form to data. Instead it accepts values at the discrete states as precise and applies interpolation to specify intermediate values.

However, DDP still must rely on an analytic function to interpolate between discrete states. The most common functions are the zero'th-order nearest neighbor and the first-order linear interpolations. Hermite and spline techniques are simply more sophisticated approaches. GDP is a DDP method that uses the Hermite interpolation to determine third or higher-order polynomials that preserve the function values and gradients at endpoints of a hyper-cube defined by discrete states. Unfortunately, even third-order polynomials can create local minima that may incorrectly be sought out by an optimization routine. Because of the order of the polynomial, at least one artificial oscillation is possible. Though oscillations will generally have low amplitudes, recursive application of an optimization routine can amplify them.

If an optimization effort fails to converge due to oscillations of the Hermite interpolation, there are a variety of approaches that we can use to achieve convergence. The simplest is, if possible, to choose an initial solution

that is closer to optimality. This may prevent or limit the growth of artificial local optima induced by oscillations. Another approach is to use known characteristics of the solution to constrain either discrete or interpolated values. For example, if an expected cost function changes monotonically over certain intervals, a problem can be constrained to specify that no values interpolated in a hyper-cube lie outside the range defined by bounds of that hyper-cube.

3. Conclusions

Incorporating autocorrelation into GDP allows us to solve more realistic and practical problems than previously possible by other DDP methods. Because GDP permits us to develop global solutions that specify optimal control decisions for every possible state of a system, we have a greater opportunity to observe general solution characteristics. This may allow us to develop general guidelines for control of other problems and to justify convenient simplifying assumptions in developing models.

The extended GDP technique presented here can be used to incorporate any information that conditions the probability distribution of stochastic inputs. In addition to autocorrelation, there is often other correlated information that we can use to improve predictions of future inputs. This information is important for determining optimal control decisions and should be included in models used to determine optimal control. The extended GDP algorithm presented here is general and allows addition of any state variable to a model that uses them to condition the distribution of stochastic variables.

V. Discussion and Concluding Remarks

A. EXTENSIONS OF THIS ANALYSIS

The process of replicating a system by an abstract model involves significant simplification. In the model presented here, I have used several simplifications that may be inappropriate for the EBMUD system.

Improved Modeling of Physical Characteristics

Snow accumulates in the Sierra Nevada Mountains during wet winter months and subsequently feeds stream flow during the spring and early summer. As a result, measurements of snowpack accumulation can greatly improve estimates of stream flow and should be included in a model of the EBMUD system used to determine optimal control. Also, the system consists of two reservoirs in the Sierras and several small terminal reservoirs. While it may be reasonable to lump the storage of some of these reservoirs in the same state variable, the appropriateness of these simplifying assumptions should be checked. Similarly, the groundwater component of the proposed system has been reduced to a "bathtub" model that may not be appropriate for this conjunctive-use system.

Addition of Economic Costs

In addition, I have not included estimates of actual economic costs arising from water rationing. Economic costs become important when balancing the costs and benefits of different objectives. Even with only a water supply objective, costs of rationing must be balanced against costs of groundwater pumping. To provide an appropriate basis for including different criteria, reasonable measures of economic costs can be developed.

Effects of Discretization

There are various aspects of dynamic programming that we should consider before accepting a solution. As pointed out by Yakowitz [1982], "One should be wary of proposed solutions until the deleterious effects of discretization have been somehow bounded and found acceptable." While GDP greatly increases the solution accuracy for a particular discretization, discretization still adds some amount of error to the solution. I should evaluate this error to ensure an appropriate discretization for the solution.

As pointed out by Johnson et al. [1993], if the general form of the expected cost function or other characteristics of the model solution is known, this knowledge may be used to select appropriate discrete values of state variables. Thus, it may also be possible to provide some general guidance on the selection of appropriate discrete values.

Other Modeling Approaches

Finally, the conjunctive use problem may be solvable by simpler methods than presented here, thus allowing the development of more realistic models. As discussed earlier, we may be able to identify appropriate functions that can be used to avoid discretization of some or all state variables. Alternatively, we may be able to break the problem into simpler sub-problems or restrict our solution to avoid considering uninteresting states or decisions.

B. CONCLUSIONS

This work has expanded the ability of dynamic programming techniques to solve a wider variety of stochastic problems by Gradient Dynamic Programming. By the expanded algorithm presented here, I have addressed a problem that contains autocorrelation without resorting to heuristic constraints on the control of groundwater.

The solution of a generic conjunctive use problem has allowed us to develop some initial insight into the optimal control of these problems. This problem will form the basis for problems that model the system more realistically. This will promote the goal of providing practical guidance to planners and managers of conjunctive use systems attempting to use a conjunctive-use approach.

The Gradient Dynamic Programming method presented here also has application beyond that of conjunctive use. It can be a useful tool for any stochastic dynamic programming problem, even beyond those of water resource management.

C. DEFINITION OF VARIABLES

$\mathbf{u}$	vector of control decision variables $u_j, j = 1, \dots, n$
$\mathbf{u}^*$	vector of optimal control decisions
$\mathbf{x}$	vector of state variables $x_j, j = 1, \dots, n$
$\mathbf{x}_{(i)}$	values of state variables at discrete state $(i), i = 1, \dots, V^n$
$\mathbf{y}$	vector of state variables at the end of the current decision period
$\mathbf{w}$	vector of stochastic variables $w_j, j = 1, \dots, n$
$\mathbf{w}_{(k)}$	values of stochastic variables at discrete realization (k)
$\mathbf{T}$	vector of transition functions such that $\mathbf{y} = \mathbf{T}(\mathbf{x}, \mathbf{u}, \mathbf{w})$
$g, \mathbf{g}$	function and coefficient vector defining the left-hand side of constraints
$\mathbf{G}$	array of left-hand side coefficients of currently active constraint set
d	function defining the right-hand side of constraints
$\mathbf{d}$	vector of right-hand side coefficients of currently active constraint set
f	total cost function, $f(\mathbf{x}, \mathbf{u}, \mathbf{w}) = C + F$
f^*	optimal total cost function, $f^* = f(\mathbf{x}, \mathbf{u}^*, \mathbf{w})$
C	current cost function, $C(\mathbf{x}, \mathbf{u})$
F	expected future cost or "cost-to-go" function, $F(\mathbf{y})$
$E_{\mathbf{w}}\{\}$	expected value operator with respect to $\mathbf{w}$
p_k	probability weight for $\mathbf{w}_{(k)}$
t	index for stage
τ	index for month
w_t	a single stochastic variable w at time t
ω_τ	log-transformed stream flow for month τ using the 3-parameter model
μ_τ	mean of ω_τ
σ_τ	standard deviation of ω_τ
χ_τ	third parameter in the three-parameter lognormal distribution
ρ_τ	lag-one autocorrelation coefficient of log-transformed stream flows
ε	random variable with normal (0,1) distribution
q	quantile of a probability distribution
∂	partial derivative operator
d	total derivative operator (defined by the chain-rule)
f_x	shorthand notation for partial derivatives (here of f with respect to $\mathbf{x}$) used with scalar functions (eg. f, C, F)

$\nabla_{\mathbf{x}} f$	shorthand notation for total derivatives (here of f with respect to $\mathbf{x}$) used with scalar functions
π_{τ}	partial autocorrelation coefficient
c	standardization constant applied to expected total cost function
N	number of stages in the planning horizon
$\mathbf{R}$	matrix of partial derivatives $\partial w / \partial \mathbf{x}$
r_{ij}	elements of $\mathbf{R}$
$\mathbf{P}$	matrix of correlation coefficients
α	discount rate
β	discount factor
V	average number of discrete values to span each state variable
n	dimension of problem (i.e., number of state variables)
T	effort to determine $\mathbf{u}^*$ for one discrete state $\mathbf{x}_{(i)}$ and stage t
Δx	discretization interval
h	total cost function augmented with constraints by Lagrange multiplier method
λ	Lagrange multipliers that minimize h
$\mathbf{D}$	shorthand notation used for a defined matrix calculation
γ	distance from solution $\mathbf{u}_0$ to the first new constraint in direction $\Delta \mathbf{u}$
ϕ	Hermite weights applied to function values
ψ	Hermite weights applied to function derivatives
$\mathbf{x}^a, \mathbf{x}^b$	upper and lower hypercube bounds
ξ_k	distance between $\mathbf{x}$ and $\mathbf{x}^b$ in dimension k
η_k	distance between $\mathbf{x}$ and $\mathbf{x}_{(i)}$ in dimension k

Appendix A

Constraining the Optimization Problem

The goal is to provide rules that allow the system operator to make real-time operating decisions. These rules determine the control decisions $\mathbf{u}^*$ that minimize an objective function f ; in other words

$$\min_{\mathbf{u}} \{f(\mathbf{x}, \mathbf{u}, \mathbf{w}, \mathbf{y})\} \quad \text{subject to active constraints} \quad \mathbf{G}\mathbf{u} = \mathbf{d}(\mathbf{x})$$

where $\mathbf{G}$ is a $(p \times m)$ matrix of rank p , and $\mathbf{d}(\mathbf{x})$ is a vector of the m right hand sides. The general form of each inequality constraint is

$$\mathbf{g}^T \mathbf{u} \leq d(\mathbf{x})$$

where $\mathbf{g}$ is a vector of coefficients and d can be a function of $\mathbf{x}$. When binding, an inequality constraint becomes an equality constraint in the active set.

An empty feasible decision space may result if the algorithm adds more than one constraint to the active set at a time. When an improvement $\Delta \mathbf{u}$ violates more than one constraint, the algorithm adds to the active set only the first constraint encountered when moving from $\mathbf{u}_0$ in the direction $\Delta \mathbf{u}$. The distance from $\mathbf{u}_0$ to each new constraint is

$$\gamma = (d - \mathbf{g}^T \mathbf{u}) / \mathbf{g}^T \Delta \mathbf{u}$$

and the first violated constraint is that with the smallest distance γ .

1. Formulation of the Constraints

A constraint on the decision variables $\mathbf{u}$ has the general form

$$g(\mathbf{u}) \leq d(\mathbf{x}, \mathbf{w})$$

where $g(\mathbf{u})$ is a function of $\mathbf{u}$ and right hand side d is a function of $\mathbf{x}$ and $\mathbf{w}$. Function $d(\mathbf{x}, \mathbf{w})$ is not dependent on $\mathbf{y}$ since $\mathbf{y}$ can be replaced by the transition function $\mathbf{y} = \mathbf{T}(\mathbf{x}, \mathbf{u}, \mathbf{w})$. The "less-than-or-equal-to" form is standard for non-linear optimization and allows consistent interpretation of the Lagrange multipliers when verifying the active constraint set. Equality constraints are always active.

Assuming a system with linear dynamics (i.e., no second order terms in the transfer function), the terms containing u can be separated from those with x and w . Thus, the constraints can be represented by

$$g^T u \leq d$$

where g is a vector of coefficients. Limited non-linear dynamics can be represented by this representation of constraints as long as non-linear terms are limited to $d(x, w)$.

The active set method used to solve the problem relies on continuously updating a set of constraints limited to those that bind the current solution. Constraints, when active, become equality constraints, and the entire active set can be represented by

$$Gu = d$$

where G is a scalar matrix of coefficients, and d is a vector of right hand.

Vector d is, most generally, a function of x and w . The dependency on w can be removed by assuming an appropriately conservative value of w . By over-constraining u , a solution will always (or practically always) be feasible, regardless of w . This is required if we are to make decision that are applied throughout a stage. If this restriction is inefficient, shorter periods should be used.

For example, suppose that w_j represents the current stage's unknown stream flow, and u_k is the current decision to supply water. If we assume w_j is zero, a manager is constrained to use only previously stored water. Thus, regardless of the actual flow realized, the manager will not make a decision to supply water that may not be available.

If the distribution of w_j is unbounded, we can still apply this approach by using a sufficiently unlikely value to bound the distribution. For example, in flood control problems it may be useful to use the 10,000 year flood value. For a normally distributed w_j , a value which has a standard deviation of greater than four often will be sufficiently conservative (probability of exceedence $< 10^{-4}$). If the control decisions are overly constrained by the resulting values of w_j , a shorter stage should be used.

2. Types of Constraints

The type of constraints used in a problem will often be a mathematical representation of physical limits of the specific system being studied. Though the constraints will depend on the specific problem, most generally they will be of two forms: (a) bounds on decision variables, $\mathbf{u}$, and (b) bounds on state variables, $\mathbf{x}$ and $\mathbf{y}$. In either case, the mathematical forms of the constraints are reorganized as above to be equality constraints that operate on $\mathbf{u}$.

(a) Bounds on Decision Variables

These will have the general form

$$u_{j,\min} \leq u_j \leq u_{j,\max}$$

where u_j is single decision of the entire set $\mathbf{u}$. When active, this constraint is represented by

$$\mathbf{g}^T \mathbf{u} = d$$

where $\mathbf{g}$ is a vector of zeros except for element k , which is either one (if $d = -u_{j,\min}$), or minus one (if $d = u_{j,\max}$). Since the right-hand side is a scalar, $dd/d\mathbf{x} = 0$.

(b) Bounds on State Variables

These will have the general form

$$x_{j,\min} \leq x_j \leq x_{j,\max} \quad \text{and} \quad y_{j,\min} \leq y_j \leq y_{j,\max}$$

where x_j and y_j are the j 'th state variables describing the current and end-of-stage states of the system. The constraints on these two state variables will often be the same since they represent the same system characteristic at adjacent stages. They can be different, such as when constraints change with time. For example, the maximum reservoir storage available for water supply may change with the season due to capacity reserved for flood control.

Constraints on $\mathbf{x}$ can be ignored since we choose the initial state of the system, and we only consider initial states that are feasible. Decisions affect $\mathbf{y}$ through the transfer function $\mathbf{y} = \mathbf{T}(\mathbf{x}, \mathbf{u}, \mathbf{w})$; and, as discussed above, constraints on $\mathbf{y}$ will be reorganized by substituting the transfer function representation

and moving all terms with $\mathbf{u}$ to the left hand side. When active, this constraint is again represented by

$$\mathbf{g}^T \mathbf{u} = d(\mathbf{x})$$

and $dd/d\mathbf{x} = \partial d/\partial \mathbf{x}$.

Appendix B

Function Approximation Through Hermite Interpolation

Given the function value of $F(\mathbf{x})$ and vector of derivatives $\partial F/\partial \mathbf{x}$ at all surrounding nodes $\mathbf{x}_i$, we can interpolate the values at intermediate states by Hermite interpolation following the method presented by *Kitanidis* [1986] and summarized by *Foufoula-Georgiou and Kitanidis* [1988, Appendix A and B]. This is accomplished by using a weighted sum of the values at the $i = 1, \dots, 2^n$ corner points of the hyper-cube that bound the intermediate point (Figure B-1). Unlike the convention of the rest of the text, here I use subscripts only as identifiers and they do not indicate derivatives.

The weightings ϕ_i and ψ_{ij} applied to node i of the hyper-cube are determined by simple polynomials that have the following properties:

	<u>0'th Order Value</u>	<u>1st Derivatives</u>
ϕ_i :	1 at node i , 0 otherwise	0 at all nodes
ψ_{ij} :	0 at all nodes and all directions	1 at node i and direction j , 0 otherwise.

Given these properties of ϕ_i and ψ_{ij} , the values of the function and its derivative can be approximated as

$$F = \sum_{i=1}^{2^n} \{F_i \phi_i + \frac{\partial F_i}{\partial \mathbf{x}} \psi_i\}$$

$$\frac{\partial F}{\partial \mathbf{x}} = \sum_{i=1}^{2^n} \{F_i \frac{\partial \phi_i}{\partial \mathbf{x}} + \frac{\partial F_i}{\partial \mathbf{x}} \frac{\partial \psi_i}{\partial \mathbf{x}}\}$$

where $\psi_i = (\psi_{i1}, \dots, \psi_{in})^T$.

The simplest polynomials that we can use to specify ϕ_i and ψ_{ij} has non-zero third derivatives and continuous second derivatives. Thus we can define a continuous second derivative of $F(\mathbf{x})$ to be

$$\frac{\partial^2 F}{\partial \mathbf{x} \partial \mathbf{x}} = \sum_{i=1}^{2^n} \{F_i \frac{\partial^2 \phi_i}{\partial \mathbf{x} \partial \mathbf{x}} + \frac{\partial F_i}{\partial \mathbf{x}} \frac{\partial^2 \psi_i}{\partial \mathbf{x} \partial \mathbf{x}}\}$$

where

$$\frac{\partial F_i}{\partial \mathbf{x}} \frac{\partial^2 \psi_i}{\partial \mathbf{x} \partial \mathbf{x}} = \sum_{j=1}^n \left\{ \frac{\partial F_i}{\partial \mathbf{x}} \frac{\partial^2 \psi_{ij}}{\partial \mathbf{x} \partial \mathbf{x}} \right\} .$$

Let $\mathbf{x}^a$ and $\mathbf{x}^b$ be vectors containing respectively the upper and lower bounds of the hyper-cube in each of the $k = 1, \dots, n$ dimensions. The dimensions of the hyper-cube are then $\Delta \mathbf{x} = \mathbf{x}^a - \mathbf{x}^b$. The location of the intermediate point can then be defined as $\xi_k = (x_k - x_k^b)/\Delta x_k$ and $0 \leq \xi_k \leq 1$.

The distance of the point from any given corner $\mathbf{x}_i$ can then be specified as $\boldsymbol{\eta} = (\eta_1, \dots, \eta_n)^T$ where

$$\begin{aligned} \eta_k &= \xi_k & \text{if } x_k &= x_k^b \\ \eta_k &= (1 - \xi_k) & \text{if } x_k &= x_k^a . \end{aligned}$$

The weighting polynomials can be defined as follows

$$\phi_i = RP$$

$$\psi_{ij} = \eta_j(1 - \eta_j)P$$

where

$$R = 1 + \sum_{k=1}^n \eta_k - 2 \sum_{k=1}^n \eta_k^2$$

$$P = \prod_{k=1}^n \{1 - \eta_k\} .$$

The above equations could also use vector notation; thus $R = 1 + (1 - 2\boldsymbol{\eta})^T \boldsymbol{\eta}$.

The first derivatives of these polynomials are then

$$\frac{d\phi_i}{dx_s} = \frac{d\eta_s}{dx_s} [(1 - 4\eta_s)P - RP_s]$$

$$\frac{\partial \psi_{ij}}{\partial x_s} = \frac{\partial \eta_s}{\partial x_s} \left(\frac{\partial \eta_j}{\partial x_j} \right)^{-1} TP_s$$

where

$$P_{a,b} = \prod_{\substack{k=1 \\ k \neq a,b}}^n \{1 - \eta_k\}$$

$$T = (1 - \eta_j)(1 - 3\eta_j) \quad \text{if } s = j$$

$$T = \eta_j(\eta_j - 1) \quad \text{if } s \neq j$$

The second derivatives of these polynomials are

$$\frac{\partial^2 \phi_i}{\partial x_l \partial x_s} = \frac{\partial \eta_l}{\partial x_l} \frac{\partial \eta_s}{\partial x_s} S_{l,s}$$

$$\text{where} \quad S_{s,s} = 6(2\eta_s - 1)P_s \quad \text{if } l = s$$

$$S_{l,s} = (4\eta_s - 1)P_s + (4\eta_l - 1)P_l + RP_{l,s} \quad \text{if } l \neq s$$

$$\text{and} \quad \frac{\partial^2 \psi_{ij}}{\partial x_l \partial x_s} = \frac{\partial \eta_l}{\partial x_l} \frac{\partial \eta_s}{\partial x_s} \left(\frac{\partial \eta_j}{\partial x_j} \right)^{-1} W_{l,s,j}$$

$$\begin{aligned} \text{where} \quad W_{j,j,j} &= (6\eta_j - 4)P_j & \text{if } l = s = j \\ W_{l,j,j} &= (3\eta_j - 1)P_l & \text{if } l \neq j, s = j \\ W_{j,s,j} &= (3\eta_j - 1)P_s & \text{if } l = j, s \neq j \\ W_{s,s,j} &= 0 & \text{if } l \neq s, s \neq j \\ W_{l,s,j} &= \eta_j(1 - \eta_j)P_{l,s} & \text{if } l \neq s, l \neq j, s \neq j \end{aligned}$$

This presentation is essentially the same as that of *Foufoula-Georgiou and Kitanidis* [1988] except for correcting equation B7.

Hypercube of 3 Dimensions

- discrete $\mathbf{x}(i) = [x_1(i), x_2(i), x_3(i)]$, for $i = 1, 2, 3, \dots, 2^3$
- intermediate $\mathbf{x} = [x_1, x_2, x_3]$
 - or $= [(x_1^b + \xi_1 \Delta x_1), (x_2^b + \xi_2 \Delta x_2), (x_3^b + \xi_3 \Delta x_3)]$
 - or $= [(x_1(1) + \eta_1(1) \Delta x_1), (x_2(1) + \eta_2(1) \Delta x_2), (x_3(1) + \eta_3(1) \Delta x_3)]$
 - or $= [(x_1(2) - \eta_1(2) \Delta x_1), (x_2(2) + \eta_2(2) \Delta x_2), (x_3(2) + \eta_3(2) \Delta x_3)]$
 - or $= [(x_1(3) + \eta_1(3) \Delta x_1), (x_2(3) - \eta_2(3) \Delta x_2), (x_3(3) + \eta_3(3) \Delta x_3)]$
 - $= \dots$
 - or $= [(x_1(8) - \eta_1(8) \Delta x_1), (x_2(8) - \eta_2(8) \Delta x_2), (x_3(8) - \eta_3(8) \Delta x_3)]$




Fig. B-1. Illustration of a Hypercube of 3 Dimensions

Appendix C

Development of Derivatives

The most general form of the objective function can be represented as any function $f = f(x, u, w, y)$ where $u^*(x)$, $w(x)$, and $y = T(x, u, w)$. Based on these functional relationships, the following relationships can be derived:

$$\frac{dy}{du} = \frac{\partial y}{\partial u}, \quad \frac{du}{dx} = \frac{\partial u}{\partial x}, \quad \frac{dw}{dx} = \frac{\partial w}{\partial x},$$

$$\frac{dy}{dx} = \frac{\partial y}{\partial x} + \frac{\partial y}{\partial u} \frac{\partial u}{\partial x} + \frac{\partial y}{\partial w} \frac{\partial w}{\partial x}.$$

In the process of developing the following derivatives, I assume that $y = T(x, u, w)$ is a linear function. This allows us to drop second-order derivatives of y and improve the presentation's clarity. If the transition function were non-linear, these terms could be preserved with only minor modification to the following presentation. To the more general formulation, I apply the objective function proposed by *Foufoula-Georgiou and Kitanidis* [1988] for reservoir management:

$$f = C(x, u) + F(y) \quad (C1)$$

where $C(x, u)$ is a single-stage cost function, and $F(y)$ is a cost function incorporating expected costs of future stages.

A. Derivatives of the Objective Function, $f(x, u, w, y)$

(1) Derivatives of objective function with respect to decisions:

$$\nabla_u f = f_u + f_y \frac{\partial y}{\partial u} \quad (C1a)$$

and, using the more specific objective function (C1),

$$\nabla_u f = C_u + F_y \frac{\partial y}{\partial u} \quad (C1b)$$

(2) Second derivatives of objective function with respect to decisions:

$$\nabla_u (\nabla_u f)^T = \frac{\partial (\nabla_u f)^T}{\partial u} + \frac{\partial (\nabla_u f)^T}{\partial y} \frac{\partial y}{\partial u}.$$

Applying the above relationship for $\nabla_{\mathbf{w}} f$ and dropping second order derivatives of $\mathbf{y}$:

$$\nabla_{\mathbf{u}}(\nabla_{\mathbf{w}} f)^T = f_{\mathbf{uu}} + \frac{\partial \mathbf{y}^T}{\partial \mathbf{u}} f_{\mathbf{yu}} + f_{\mathbf{uy}} \frac{\partial \mathbf{y}}{\partial \mathbf{u}} + \frac{\partial \mathbf{y}^T}{\partial \mathbf{u}} f_{\mathbf{yy}} \frac{\partial \mathbf{y}}{\partial \mathbf{u}}, \quad (\text{C2a})$$

and, more specifically, $\nabla_{\mathbf{u}}(\nabla_{\mathbf{w}} f)^T = C_{\mathbf{uu}} + \frac{\partial \mathbf{y}^T}{\partial \mathbf{u}} F_{\mathbf{yy}} \frac{\partial \mathbf{y}}{\partial \mathbf{u}}.$ (C2b)

(3) Derivatives of objective function with respect to initial state:

$$\nabla_{\mathbf{x}} f = f_{\mathbf{x}} + f_{\mathbf{u}} \frac{\partial \mathbf{u}}{\partial \mathbf{x}} + f_{\mathbf{y}} \frac{d\mathbf{y}}{d\mathbf{x}} + f_{\mathbf{w}} \frac{\partial \mathbf{w}}{\partial \mathbf{x}}.$$

Expanding $d\mathbf{y}/d\mathbf{x}$, we can re-express this in terms of $\nabla_{\mathbf{w}} f$ as follows

$$\nabla_{\mathbf{x}} f = \nabla_{\mathbf{u}} f \frac{\partial \mathbf{u}}{\partial \mathbf{x}} + f_{\mathbf{x}} + f_{\mathbf{y}} \left(\frac{\partial \mathbf{y}}{\partial \mathbf{x}} + \frac{\partial \mathbf{y}}{\partial \mathbf{w}} \frac{\partial \mathbf{w}}{\partial \mathbf{x}} \right) + f_{\mathbf{w}} \frac{\partial \mathbf{w}}{\partial \mathbf{x}} \quad (\text{C3a})$$

and $\nabla_{\mathbf{x}} f = \nabla_{\mathbf{u}} f \frac{\partial \mathbf{u}}{\partial \mathbf{x}} + C_{\mathbf{x}} + F_{\mathbf{y}} \left(\frac{\partial \mathbf{y}}{\partial \mathbf{x}} + \frac{\partial \mathbf{y}}{\partial \mathbf{w}} \frac{\partial \mathbf{w}}{\partial \mathbf{x}} \right).$ (C3b)

(4) Derivatives of objective function with respect to initial state and decisions

(Note that $\nabla_{\mathbf{u}}(\nabla_{\mathbf{x}} f)^T = (\nabla_{\mathbf{x}}(\nabla_{\mathbf{u}} f)^T)^T$):

$$\nabla_{\mathbf{u}}(\nabla_{\mathbf{x}} f)^T = \left(\frac{\partial(\nabla_{\mathbf{u}} f)^T}{\partial \mathbf{x}} + \frac{\partial(\nabla_{\mathbf{u}} f)^T}{\partial \mathbf{u}} \frac{\partial \mathbf{u}}{\partial \mathbf{x}} + \frac{\partial(\nabla_{\mathbf{u}} f)^T}{\partial \mathbf{w}} \frac{\partial \mathbf{w}}{\partial \mathbf{x}} + \frac{\partial(\nabla_{\mathbf{u}} f)^T}{\partial \mathbf{y}} \frac{d\mathbf{y}}{d\mathbf{x}} \right)^T.$$

Applying the relationship for $\nabla_{\mathbf{w}} f$ and dropping second order derivatives of $\mathbf{y}$:

$$\begin{aligned} \nabla_{\mathbf{u}}(\nabla_{\mathbf{x}} f)^T &= f_{\mathbf{xu}} + f_{\mathbf{xy}} \frac{\partial \mathbf{y}}{\partial \mathbf{u}} + \frac{\partial \mathbf{u}^T}{\partial \mathbf{x}} \left(f_{\mathbf{uu}} + f_{\mathbf{uy}} \frac{\partial \mathbf{y}}{\partial \mathbf{u}} \right) + \frac{\partial \mathbf{w}^T}{\partial \mathbf{x}} \left(f_{\mathbf{wu}} + f_{\mathbf{wy}} \frac{\partial \mathbf{y}}{\partial \mathbf{u}} \right) \\ &\quad + \frac{d\mathbf{y}^T}{d\mathbf{x}} \left(f_{\mathbf{yu}} + f_{\mathbf{yy}} \frac{\partial \mathbf{y}}{\partial \mathbf{u}} \right) \end{aligned}$$

Expanding $d\mathbf{y}/d\mathbf{x}$, we can re-express this in terms of $\nabla_{\mathbf{u}}(\nabla_{\mathbf{w}} f)^T$ as follows

$$\begin{aligned} \nabla_{\mathbf{u}}(\nabla_{\mathbf{x}} f)^T &= \frac{\partial \mathbf{u}^T}{\partial \mathbf{x}} \nabla_{\mathbf{u}}(\nabla_{\mathbf{w}} f)^T + f_{\mathbf{xu}} + f_{\mathbf{xy}} \frac{\partial \mathbf{y}}{\partial \mathbf{u}} + \frac{\partial \mathbf{w}^T}{\partial \mathbf{x}} \left(f_{\mathbf{wu}} + f_{\mathbf{wy}} \frac{\partial \mathbf{y}}{\partial \mathbf{u}} \right) \\ &\quad + \left(\frac{\partial \mathbf{y}}{\partial \mathbf{x}} + \frac{\partial \mathbf{y}}{\partial \mathbf{w}} \frac{\partial \mathbf{w}}{\partial \mathbf{x}} \right)^T \left(f_{\mathbf{yu}} + f_{\mathbf{yy}} \frac{\partial \mathbf{y}}{\partial \mathbf{u}} \right) \end{aligned} \quad (\text{C4a})$$

$$\text{and} \quad \nabla_u(\nabla_x f)^T = \frac{\partial u^T}{\partial x} \nabla_u(\nabla_w f)^T + C_{xu} + \left(\frac{\partial y}{\partial x} + \frac{\partial y}{\partial w} \frac{\partial w}{\partial x} \right)^T F_{yy} \frac{\partial y}{\partial u} . \quad (\text{C4b})$$

The above form is useful since terms with $\partial u/\partial x$ are aggregated.

Alternately, (C3a) and (C3b) can be expressed as

$$\nabla_x f = C_x + C_u \frac{\partial u}{\partial x} + F_y \frac{dy}{dx} \quad (\text{C5a})$$

$$\nabla_u(\nabla_x f)^T = C_{xu} + \frac{\partial u^T}{\partial x} C_{uu} + \frac{dy^T}{dx} F_{yy} \frac{\partial y}{\partial u} , \quad (\text{C5b})$$

but these forms are not used except for comparison with the work of *Foufoula-Georgiou and Kitanidis* [1988].

MODIFICATION OF PRIOR WORK

Foufoula-Georgiou and Kitanidis [1988] proposed the following derivations

$$\nabla_w f = \nabla_u C + \nabla_y F \frac{\partial y}{\partial u} , \quad (\text{C16})$$

$$f_{uu} = C_{uu} + \frac{\partial y^T}{\partial u} F_{yy} \frac{\partial y}{\partial u} , \quad (\text{C17})$$

$$\nabla_x f = \nabla_x C + \nabla_u C \frac{\partial u}{\partial x} + \nabla_y F \left(\frac{\partial y}{\partial x} + \frac{\partial y}{\partial u} \frac{\partial u}{\partial x} \right) , \quad (\text{C18})$$

$$\nabla_x(\nabla_w f)^T = C_{ux} + \frac{\partial y^T}{\partial u} F_{yy} . \quad (\text{C19})$$

It is useful to first note the following relationships:

$$\nabla_u C = C_u , \quad \nabla_y F = F_y , \quad \nabla_y(\nabla_y F)^T = F_{yy} ,$$

$$\nabla_x C = C_x + C_u \frac{\partial u}{\partial x} , \text{ and } \nabla_x(\nabla_u C)^T = C_{ux} + C_{uu} \frac{\partial u}{\partial x} .$$

The following observations are made in comparing the results of the derivations presented here and those of the original work:

(1) Equation (C16) is the same because of the equivalence of the general and partial derivatives.

(2) Equation (C17) is also the same except for the inconsistent use of the symbol f_{uu} which implies a partial derivative, when the general derivative is desired.

(3) The first two terms of equation (C18) are different because of the incorrect use of the general derivative symbols ∇ and d when partial derivatives should have been used:

$$\nabla_x C = C_x + C_u \frac{\partial u}{\partial x} \neq \nabla_x C + \nabla_u C \frac{du}{dx} .$$

(4) The definition for $\nabla_u(\nabla_x f)^T$ (equation C19) in the original work is absent the final term dy/dx . Also the first two terms are different because of the incorrect use of the partial derivative notation C_{ux} when the general derivative should have been specified:

$$\nabla_x(\nabla_u C)^T = C_{ux} + C_{uu} \frac{\partial u}{\partial x} \neq C_{ux}$$

(5) The total derivative

$$\frac{dy}{dx} = \frac{\partial y}{\partial x} + \frac{\partial y}{\partial u} \frac{\partial u}{\partial x} + \frac{\partial y}{\partial w} \frac{\partial w}{\partial x}$$

presented here is consistent with the original work since, in the absence of autocorrelation, $\partial w/\partial x = 0$; thus

$$\frac{dy}{dx} = \frac{\partial y}{\partial x} + \frac{\partial y}{\partial u} \frac{\partial u}{\partial x} .$$

When modeling a system with autocorrelation, we need the longer form to evaluate $\nabla_x f$ and $\nabla_u(\nabla_x f)^T$.

(6) Because of autocorrelation, the matrix dy/dx is not constant, and must be recalculated for each $x_{(i)}$ and for each $w_{(k)}$. Since

$$\frac{dy_{(k)}}{dx} = \frac{\partial y}{\partial x} + \frac{\partial y}{\partial u} \frac{\partial u}{\partial x} + \frac{\partial y_{(k)}}{\partial w_{(k)}} \frac{\partial w_{(k)}}{\partial x} ,$$

each stochastic realization and discrete state affects the calculated value through the additional terms in this formulation. Even with a linear transition function, the above total derivative is not constant because of the effect that state variables have on conditioning the distribution of w .

B. Approximation by Discrete Realizations

The values derived from the above formulas are uncertain because of the presence of stochastic variables $\mathbf{w}$. We can estimate these values by using probability weighted realizations $\mathbf{w}_{(k)}$ to approximate their expected values. We can choose discrete realizations $\mathbf{w}_{(k)}$ to span the multi-variate distribution of $\mathbf{w}$ much as we can choose states $\mathbf{x}_{(i)}$ to span $\mathbf{x}$. The objective function (C1) is evaluated as

$$f = C(\mathbf{x}, \mathbf{u}) + \sum_k \{p_k F(\mathbf{y}_{(k)})\} \quad (\text{C6})$$

where p_k are the weights and $\mathbf{y}_{(k)} = \mathbf{T}(\mathbf{x}, \mathbf{u}, \mathbf{w}_{(k)})$. Similarly,

$$\nabla_{\mathbf{u}} f = C_{\mathbf{u}} + \sum_k \{p_k F_{\mathbf{y},(k)}\} \frac{\partial \mathbf{y}}{\partial \mathbf{u}} \quad (\text{C7})$$

$$\nabla_{\mathbf{u}} (\nabla_{\mathbf{u}} f)^T = C_{\mathbf{u}\mathbf{u}} + \frac{\partial \mathbf{y}}{\partial \mathbf{u}}^T \sum_k \{p_k F_{\mathbf{y}\mathbf{y},(k)}\} \frac{\partial \mathbf{y}}{\partial \mathbf{u}} \quad (\text{C8})$$

$$\nabla_{\mathbf{x}} f = \nabla_{\mathbf{u}} f \frac{\partial \mathbf{u}}{\partial \mathbf{x}} + C_{\mathbf{x}} + \sum_k \{p_k F_{\mathbf{y},(k)}\} \frac{\partial \mathbf{y}}{\partial \mathbf{x}} + \sum_k \{p_k F_{\mathbf{y},(k)}\} \frac{\partial \mathbf{y}_{(k)}}{\partial \mathbf{w}_{(k)}} \frac{\partial \mathbf{w}_{(k)}}{\partial \mathbf{x}} \quad (\text{C9})$$

$$\begin{aligned} \nabla_{\mathbf{u}} (\nabla_{\mathbf{x}} f)^T &= \frac{\partial \mathbf{u}}{\partial \mathbf{x}}^T \nabla_{\mathbf{u}} (\nabla_{\mathbf{u}} f)^T + C_{\mathbf{x}\mathbf{u}} + \frac{\partial \mathbf{y}}{\partial \mathbf{x}}^T \sum_k \{p_k F_{\mathbf{y}\mathbf{y},(k)}\} \frac{\partial \mathbf{y}}{\partial \mathbf{u}} \\ &+ \sum_k \{p_k \frac{\partial \mathbf{w}_{(k)}}{\partial \mathbf{x}}^T \frac{\partial \mathbf{y}_{(k)}}{\partial \mathbf{w}_{(k)}}^T F_{\mathbf{y}\mathbf{y},(k)}\} \frac{\partial \mathbf{y}}{\partial \mathbf{u}} \end{aligned} \quad (\text{C10})$$

We can calculate values for equations (C6) - (C10) quickly once we determine

$$\sum_k \{p_k F_{(k)}\} , \quad (\text{C11})$$

$$\sum_k \{p_k F_{\mathbf{y},(k)}\} , \quad (\text{C12})$$

$$\sum_k \{p_k F_{\mathbf{y}\mathbf{y},(k)}\} , \quad (\text{C13})$$

$$\sum_k \{p_k F_{\mathbf{y},(k)}\} \frac{\partial \mathbf{y}_{(k)}}{\partial \mathbf{w}_{(k)}} \frac{\partial \mathbf{w}_{(k)}}{\partial \mathbf{x}} , \quad (\text{C14})$$

and

$$\sum_k \{p_k F_{\mathbf{y}\mathbf{y},(k)}\} \frac{\partial \mathbf{y}_{(k)}}{\partial \mathbf{w}_{(k)}} \frac{\partial \mathbf{w}_{(k)}}{\partial \mathbf{x}} . \quad (\text{C15})$$

C. Solving Gradients du^*/dx

Derivatives of the Lagrangian formed from the objective function and constraints yield

$$\nabla_w f + G^T \lambda = 0$$

and

$$Gu - d = 0$$

where both u^* and λ are implicit functions of x . Derivatives with respect to x are

$$\nabla_u(\nabla_x f)^T + G^T \frac{d\lambda}{dx} = 0 ,$$

and

$$G \frac{du^*}{dx} - \frac{dd}{dx} = 0$$

where the term $\nabla_u(\nabla_x f)^T$ is defined in Appendix D. This term can be divided into two parts

$$\nabla_u(\nabla_x f)^T = \frac{\partial u^T}{\partial x} \nabla_u(\nabla_w f)^T + D$$

$$\text{with } D = f_{xu} + f_{xy} \frac{\partial y}{\partial u} + \frac{\partial w^T}{\partial x} \left(f_{wu} + f_{wy} \frac{\partial y}{\partial u} \right) + \left(\frac{\partial y}{\partial x} + \frac{\partial y}{\partial w} \frac{\partial w}{\partial x} \right)^T \left(f_{yu} + f_{yy} \frac{\partial y}{\partial u} \right) .$$

This form has the advantage that terms with du^*/dx are grouped and permits us to solve du^*/dx by the simultaneous equations

$$\begin{bmatrix} \nabla_u(\nabla_w f)^T & G^T \\ G & 0 \end{bmatrix} \begin{bmatrix} du^*/dx \\ d\lambda/dx \end{bmatrix} = \begin{bmatrix} -D \\ \nabla_x d \end{bmatrix}$$

where D is defined as above, and $\nabla_x d = d_x$.

Using the objective function proposed by *Foufoula-Georgiou and Kitanidis* [1988],

$$D = C_{xu} + \left(\frac{\partial y}{\partial x} + \frac{\partial y}{\partial w} \frac{\partial w}{\partial x} \right)^T F_{yy} \frac{\partial y}{\partial u} .$$

Because D contains stochastic elements, its expected value is the probability weighted estimate

$$\mathbf{D} = \mathbf{C}_{xu} + \frac{\partial \mathbf{y}^T}{\partial \mathbf{x}} \sum_k \{p_k F_{yy,(k)}\} \frac{\partial \mathbf{y}}{\partial \mathbf{u}} + \sum_k \{p_k \frac{\partial \mathbf{w}_{(k)}}{\partial \mathbf{x}} \frac{\partial \mathbf{y}_{(k)}}{\partial \mathbf{w}_{(k)}}^T F_{yy,(k)}\} \frac{\partial \mathbf{y}}{\partial \mathbf{u}}$$

where the terms involved in summation are the same as calculated in Appendix D.

MODIFICATION OF PRIOR WORK

The original work of *Foufoula-Georgiou and Kitanidis* [1988] specified that, in a deterministic problem, the upper right-hand-side sub-matrix $[-\mathbf{D}]$ was

$$-\left(\mathbf{C}_{ux} + \frac{\partial \mathbf{y}^T}{\partial \mathbf{u}} \mathbf{F}_{yy} \right).$$

The transpose of the matrix defined by this formula yields nearly the same result as the present formulation in the absence of autocorrelation. Without state variables that incorporate autocorrelation, $\partial \mathbf{w}/\partial \mathbf{x}$ is a zero matrix and the term $\partial \mathbf{y}/\partial \mathbf{x}$ will often be an identity matrix.

Foufoula-Georgiou and Kitanidis [1988] refer to the above formula as $-\nabla_{\mathbf{x}}(\nabla_{\mathbf{u}})^T$ which is incorrect. The correct value can be found above.

References

- Bogle, M. G., and M. J. O'Sullivan, Stochastic optimization of a water supply system, *Water Resour. Res.*, 15(4), 778-786, 1979.
- Box, G. E. P., and G. M. Jenkins, *Time Series Analysis, forecasting and control*, rev. ed., Holden-Day, Oakland, Calif., 1976.
- Bras, R. L., R. Buchanan, and K. C. Curry, Real time adaptive closed loop control of reservoirs with the High Aswan Dam as a case study, *Water Resour. Res.*, 19(1), 33-52, 1983.
- Buras, N., Conjunctive operation of dams and aquifers, *J. Hydraulics Div. Amer. Soc. Civ. Eng.*, 89(HY6), 111-131, 1963.
- Buras, N., *Scientific Allocation of Water Resources*, Elsevier, New York, 1972.
- Burges, S. J., and R. Maknoon, *A Systematic Examination of Issues in Conjunctive Use of Ground and Surface Waters*, Technical Report No. 44, Charles W. Harris Hydraulics Lab., Univ. of Washington, Dept. of Civil Eng., Seattle, Wash., September, 1975.
- Cardwell, H., and H. Ellis, Stochastic dynamic programming models for water quality management, *Water Resour. Res.*, 29(4), 803-813, 1993.
- Casola, W. H., R. Narayanan, C. Duffy, and A. B. Bishop, Optimal control model for groundwater management, *J. Water Resources Plan. Manag. Amer. Soc. Civ. Eng.*, 112(2), 183-186, 1986.
- Chatfield, C., *The Analysis of Time Series, and introduction*, 4th ed., Chapman and Hall, London, 1989.
- Culver, T. B., and C. A. Shoemaker, Optimal Control for Groundwater Remediation by Differential Dynamic Programming With Quasi-Newton Approximations, *Water Resour. Res.*, 29(4), 823-831, 1993.
- Draper, N. R., and H. Smith, *Applied Regression Analysis*, 2nd ed., John Wiley & Sons, New York, 1981.
- DWR, *San Joaquin County Ground Water Investigation*, Department of Water Resources, The Resources Agency, State of California, Bulletin No. 146, July 1967.
- EBMUD, *Draft Environmental Impact Statement/Report, Updated Water Supply Management Program*, East Bay Municipal Utility District, Oakland, Calif., prepared by EDAW, Inc., San Francisco, 1992.

- Esmaeil-Beik, S., and Y.-S. Yu, Optimal operation of multipurpose pool of Elk City Lake, *J. Water Resour. Plan. Manag. Div. Am. Soc. Civ. Eng.*, 110(WR1), 1-14, 1984.
- Foufoula-Georgiou, E., and P. K. Kitanidis, Gradient dynamic programming for stochastic optimal control of multidimensional water resources systems, *Water Resour. Res.*, 24(8), 1345-1359, 1988.
- Gal, S., Optimal management of a multireservoir water supply system, *Water Resour. Res.*, 15(4), 737-749, 1979.
- Georgakakos, A. P., and D. A. Vlatas, Stochastic control of groundwater systems, *Water Resour. Res.*, 27(8), 2077-2090, 1991.
- Georgakakos, A. P., and H. Yao, New control concepts for uncertain water resources systems, 1. theory, *Water Resour. Res.*, 29(6), 1505-1516, 1993.
- Gorelick, S. M., Sensitivity analysis of optimal groundwater contaminant capture curves: Spatial variability and robust solutions, in *Solving Groundwater Problems With Models*, National Water Well Association, Denver, Colo., 1987.
- Gorelick, S. M., A review of distributed parameter groundwater management modelling methods, *Water Resour. Res.*, 19(2), 305-319, 1983.
- Hoshi, K., and S. J. Burges, Seasonal runoff volumes conditioned on forecasted total runoff volume, *Water Resour. Res.*, 16(6), 1079-1084, 1980.
- Johnson, S. A., J. R. Stedinger, and C. A. Shoemaker, Computational improvements in dynamic programming, *Forefronts* 4(7), pp. 3-7, Cent. for Theory and Simulation Sci. and Eng., Cornell Univ., 1988.
- Johnson, S. A., J. R. Stedinger, and K. Staschus, Heuristic operating policies for reservoir system simulation, *Water Resour. Res.*, 27(5), 673-685, 1991.
- Johnson, S. A., J. R. Stedinger, C. A. Shoemaker, Y. Li, and J. A. Tejada-Guibert, Numerical solution of continuous-state dynamic programs using linear and spline interpolation, *Operations Research*, submitted, 1993.
- Jones, L., R. Willis, and W. W.-G. Yeh, Optimal Control of Nonlinear Groundwater Hydraulics Using Differential Dynamic Programming, *Water Resour. Res.*, 23(11), 2097-2106, 1987.
- Karamouz, M., and H. V. Vasiliadis, Bayesian stochastic optimization of reservoir operation using uncertain forecasts, *Water Resour. Res.*, 28(5), 1221-1232, 1992.

- Karamouz, M., and M. H. Houck, Annual and monthly reservoir operating rules generated by deterministic optimization, *Water Resour. Res.*, 18(5), 1337-1344, 1982.
- Kelman, J., J. R. Stedinger, L. A. Cooper, E. Hsu, and S. Yuan, Sampling stochastic dynamic programming applied to reservoir operation, *Water Resour. Res.*, 26(3), 447-454, 1990.
- Kitanidis, P. K., and E. Foufoula-Georgiou, Error analysis of conventional discrete and gradient dynamic programming, *Water Resour. Res.*, 23(5), 845-856, 1987.
- Kitanidis, P. K., Hermite Interpolation on an n -dimensional rectangular grid, water resources technical note, St. Anthony Falls Hydraul. Lab., Univ. of Minn., Minneapolis, July 1986.
- Lettenmaier, D. P., and S. J. Burges, *Reliability of Cyclic Surface and Groundwater Storage Systems for Water Supply: A Preliminary Assessment*, Technical Report No. 64, Charles W. Harris Hydraulics Lab., Univ. of Washington, Dept. of Civil Eng., Seattle, Wash., November, 1979.
- Loucks, D. P., J. R. Stedinger, and H. A. Haith, *Water Resource Systems Planning and Analysis*, Prentice-Hall, Englewood Cliffs, N. J., 1981.
- Luenberger, D. G., *Linear and Nonlinear Programming*, 2nd ed., Addison-Wesley, Reading, Mass., 1984.
- Maddock, T., III, The operation of a stream-aquifer system under stochastic demands, *Water Resour. Res.*, 10(1), 1-10, 1974.
- McClurg, S., *Unresolved Issues in Water Marketing*, Western Water, Water Education Foundation, Sacramento, Calif., 4-11, May/June 1992b.
- McClurg, S., *Urban Water Costs*, Western Water, Water Education Foundation, Sacramento, Calif., 4-11, March/April 1992a.
- Saad, M., A. Turgeon, and J. R. Stedinger, Censored-data correlation and principal component dynamic programming, *Water Resour. Res.*, 28(8), 2135-2140, 1992.
- Saad, M., and A. Turgeon, Application of principal component analysis to long-term reservoir management, *Water Resour. Res.*, 24(7), 907-912, 1988.
- Sax, J. L., R. H. Abrams, and B. H. Thompson, *Legal Control of Water Resources*, 2nd ed., West Publishing, St. Paul, Minn., 1991.

- Stedinger, J. R. , B. R. Sule, and D. P. Loucks, Stochastic dynamic programming models for reservoir operation optimization, *Water Resour. Res.*, 20(11), 1499-1505, 1984.
- Stedinger, J. R., Estimating correlations in multivariate streamflow models, *Water Resour. Res.*, 17(1), 200-208, 1981.
- Stedinger, J. R., Fitting log normal distributions to hydrologic data, *Water Resour. Res.*, 16(3), 481-490, 1980.
- Tiedeman, C., and S. M. Gorelick, Analysis of uncertainty in optimal groundwater contaminant capture design, *Water Resour. Res.*, 29(7), 2139-2153, 1993.
- Trezos, T., and W. W-G. Yeh, Use of stochastic dynamic programming for reservoir management, *Water Resour. Res.*, 23(6), 983-996, 1987.
- Vogel, R. M., and Neil M. Fennessey, L moment diagrams should replace product moment diagrams, *Water Resour. Res.*, 29(6), 1745-1752, 1993.
- Wagner, B. J., and S. M. Gorelick, Optimal groundwater quality management under parameter uncertainty, *Water Resour. Res.*, 23(7), 1162-1174, 1987.
- Wagner, B. J., and S. M. Gorelick, Reliable aquifer remediation in the presence of spatially variable hydraulic conductivity: From data to design, *Water Resour. Res.*, 25(10), 2211-2225, 1989.
- Wagner, J. M., U. Shamir, and H. R. Nemati, Groundwater quality management under uncertainty: Stochastic programming approaches and the value of information, *Water Resour. Res.*, 28(5), 1233-1246, 1992.
- Western Water, Water Education Foundation, Sacramento, Calif., July / August 1993.
- Whittington, D., D. T. Lauria, A. M. Wright, K. Choe, J. A. Hughes, and V. Swarna, Household demand for improved sanitation services in Kumasi, Ghana: A contingent valuation study, *Water Resour. Res.*, 29(6), 1539-1560, 1993.
- Willis, R., and W. W. Yeh, *Groundwater Systems Planning & Management*, Prentice-Hall, Englewood Cliffs, N. J., 1987.
- Yakowitz, S., Dynamic programming applications in water resources, *Water Resour. Res.*, 18(4), 673-696, 1982.
- Yao, H., and A. P. Georgakakos, New control concepts for uncertain water resources systems, 2. reservoir management, *Water Resour. Res.*, 29(6), 1517-1525, 1993.

- Yazicigil, H., and M. Rasheeduddin, Optimization model for groundwater management in multi-aquifer systems, *J. Water Resour. Plan. Manage. Am. Soc. Civ. Eng.*, 113(2), 257-273, 1987.
- Yeh, W. W-G., Reservoir management and operations models: A state-of-the-art review, *Water Resour. Res.*, 21(12), 1797-1818, 1985.
- Young, G. K., Jr., Finding reservoir operating rules, *J. Hydraulics Div. Amer. Soc. Civ. Eng.*, 93(HY6), 297-321, 1967.