Methodology for Flow and Salinity Estimates in the Sacramento-San Joaquin Delta and Suisun Marsh

24th Annual Progress Report
June 2003

Chapter 2:
REALM

Author: Eli Ateljevich and Ralph Finch

2 REALM

2.1 Introduction

REALM - River, Estuary, and Land Model - is a project to develop a new modeling tool and decision support system for managing *short-term operations* in the Sacramento-San Joaquin Delta and its tributaries, and planning *long-term structural changes* such as a Through Delta Facility (TDF) or In-Delta Storage.

The REALM project is motivated by two concerns:

- ☐ The need to significantly improve the accuracy and speed of existing simulation models.
- ☐ The need to greatly increase the usefulness of models for decision-making and for solving complex water issues in the Delta, tributaries, and bays.

Like other models used in the water resources community such as DSM2, REALM will provide software for modeling physical, chemical, and biological systems in estuaries, rivers, bays, and eventually groundwater. The range of processes that can presently be modeled will be expanded, physical detail will be refined, and multi-dimensionality will be offered. Many of the simplifying assumptions that are currently made will be eliminated or made optional.

Besides better physics representations and numerical methods, REALM will offer improved support for analysis and decision-making. The user will be provided with ways to visualize and investigate tradeoffs between disparate objectives, implement adaptive operation strategies, fold field data into the model in real time, and “play” with the model interactively. The visualization and analysis tools will leverage experiments, providing information such as sensitivities and probabilities that will allowing modelers to steer their investigations towards promising operating strategies or designs.

REALM will become a vehicle for directly posing such scientific management decisions as:

- ☐ What is the most efficient flow schedule to meet the water quality standard at Rock Slough?
- ☐ How should island flooding be manipulated to obtain the best improvement of water quality?
- ☐ What is the error in this simulation due to uncertainty in the tide? What strategy best compensates for this uncertainty?
- ☐ What are the tradeoffs between pumping rate and scheduling, Clifton Court gate operations, salinity, and stage in the South Delta?

Current models merely provide the ability to simulate; and questions like the ones above can be answered only indirectly, using informal trial-and-error experiments with simplified (flat-lined) time series as input and *ad hoc* safety margins to deal with uncertainty. By augmenting the traditional water resources model with standard tools from disciplines such as optimal control, mechanical engineering, and statistics, REALM will be able to show users much more than just the simulated output associated with one choice of input. The user will learn in detail how the Bay-Delta will respond to variations in management decisions and how to better craft these decisions. Thus, REALM supports a broad range of Bay-Delta study, whether motivated by science and the need to learn about the estuary, or by practice and the need to manipulate the Delta as an engineered system.

Many problems of interest in the Delta are complex. To solve them with a usable level of accuracy requires considerable model resolution which comes at the expense of increased model run time. In addition, many of the exploratory techniques for decision support advocated in this document involve batches of related exploratory runs, not just a single simulation. As the output of analysis expands from the “single best shot” to a depiction of tradeoffs and sensitivities, computational costs will go up proportionally.

REALM addresses the need for efficiency in two ways. First, it allows adjustable or selectable computational density within a model grid. Mesh refinement is concentrated in the areas that need it most, including complex flow regions and the featured regions of the study. Second, REALM’s design allows parallel processing if multiple processor units are available in the hardware. Each processor works on a separate part of the domain or problem, with sub-components arranged in such a way that the communication between them is minimal and efficient.

2.2 REALM Physical Modeling Requirements

REALM is intended to be a state-of-the-art estuary model. The model will simulate flow, water surface height, water quality, and particle movement in response to tides, upstream flows, and midstream diversions. In its first release, REALM will include hydrodynamics and water quality in one and two dimensions plus some important extensions such as wetting and drying of tidal mudflats, and non-conservative/reactive source terms for dissolved constituents. In addition, due to strong demand in the user community, REALM will, early on, also include accurate individual particle tracking with behavior.

In the follow-up releases of REALM, vertically layered flow and 3D water quality modeling, groundwater, and sediment transport will be introduced. Care will be taken to ensure that the 3D and 2D formulations are compatible, so that coupling between them is simple and physical and the two can coexist in the same simulation (an example of this is TRIM3D, whose representation with one vertical layer is equivalent to TRIM2D.)

The order in which physical processes will be added to REALM depends on user needs. The Delta Modeling Section has interviewed more than a dozen Delta workers, managers, and stakeholders to determine their priorities. These priorities and associated development costs will continue to be assessed and balanced in a regular design cycle.

2.3 Numerical Methods and Grid

Like all estuary models, REALM will be based on conservation of mass and momentum in one, two, and three dimensions. The formulation is not discussed in any detail here, but the formulation and level of approximation will be comparable to that used in comparable international (MIKE, Telemac) and local (TRIM, RMA, DSM2) flow models. The portrayal of the physics will include hydrostatic assumption for pressure variations in the vertical direction, variable water density due to salinity (using the Boussinesq approximation), diffusion analogies for turbulent stresses and dispersion, and a variety of source and sink terms such as wind and friction stresses.

Numerical models are distinguished not only by their specification of the conservation laws but also by the numerical methods that are used to integrate the resulting equations. The REALM design will accommodate a variety of numerical schemes[1], but primarily targets Finite Volume Methods (FVM) on Cartesian cut cell grids. In FVM schemes, the study domain is divided into discrete computational cells or "volumes" and the exchange of mass and momentum between these cell is tracked over time. A great variety of finite volume methods arise from the methods for calculating intercell fluxes, the treatment of boundary conditions, the choice of grid used to represent the estuary, and the marching scheme used to advance the model in time.

2.3.1 Grid

One of the most significant choices for a FVM numerical model design is the choice of the spatial discretization or grid. The grid not only affects the accuracy, speed, and user-friendliness of the model; it also imposes a number of restrictions on the choices of specific numerical schemes and has repercussions on all aspects of the software design including data management and input/output. For one-dimensional modeling, REALM will continue to use a network of channels similar to the current DWR model DSM2. For 2D and 3D, REALM will use a hybrid mesh that combines several tools often associated with one another: multiblock rectangular meshes, Cartesian cut cells, and mesh refinement. These methods retain the excellent numerical properties that rectangular grids possess, but allow boundary-conforming geometry and concentration of computation in difficult or interesting regions.

Multiblock refers to the use of a number of rectangles instead of a single rectangle to cover the domain (Figure 2.1). "Cartesian cut cells" refers to the use of a regular mesh for most of the domain, with the irregular boundary embedded using a series of straight-line cuts (Figure 2.2). The description of the cuts can be complicated when the definition of the model domain changes with fluctuating water levels during wetting and drying (Figure 2.3). Adaptive Mesh Refinement (AMR) refers to an adaptive solution mechanism by which Cartesian grids of different levels of

---

[1] If any broad category of numerical formulation is excluded, it is finite element methods (FEM). Although reputable models such as RMA and Telemac use finite elements, these schemes are seldom used in new flow models. There is little generic software support for FEM in infrastructure libraries, and most of the innovations from FEM (including arbitrary grids) have been absorbed into the finite difference and volume methods.

refinement are nested to produce a grid with high resolution in some areas and lower resolution in others (Figure 2.4).

Cartesian cut cell and AMR methods are difficult to program but they are easy to use. To allow efficient model development, "infrastructure libraries" such as the Advanced Computational Testing and Simulation (ACTS) collection developed at the Department of Energy national laboratories, particularly LBL, will be used. LBL is expected to be a major collaborator on this project.

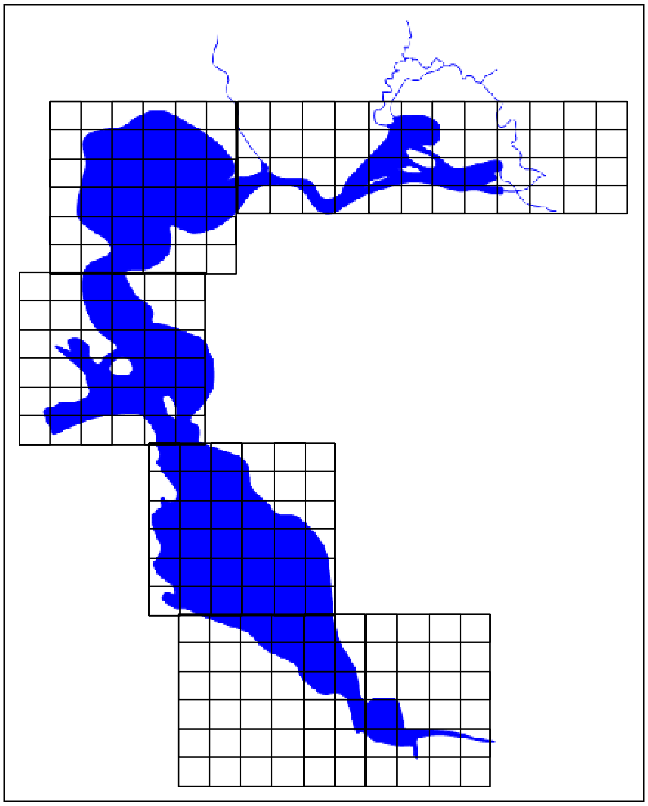

Figure 2.1: Multiblock Using Several Rectangular Meshes to Cover a Model Domain.

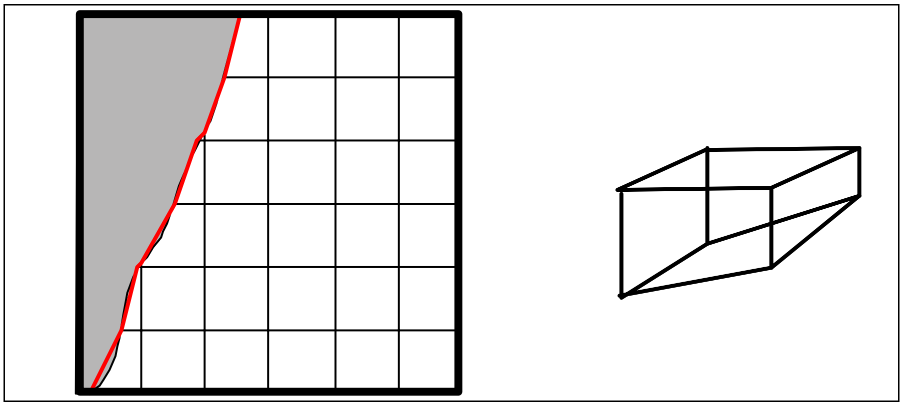

Figure 2.2: (left) Plan View of 2D Cartesian Cut Cells, with Boundary-Conforming Straight Line Cuts Across Cells. (right) Side View of a Single 3D Cut Cell, with a Planar Cut Representing Bottom Topography.

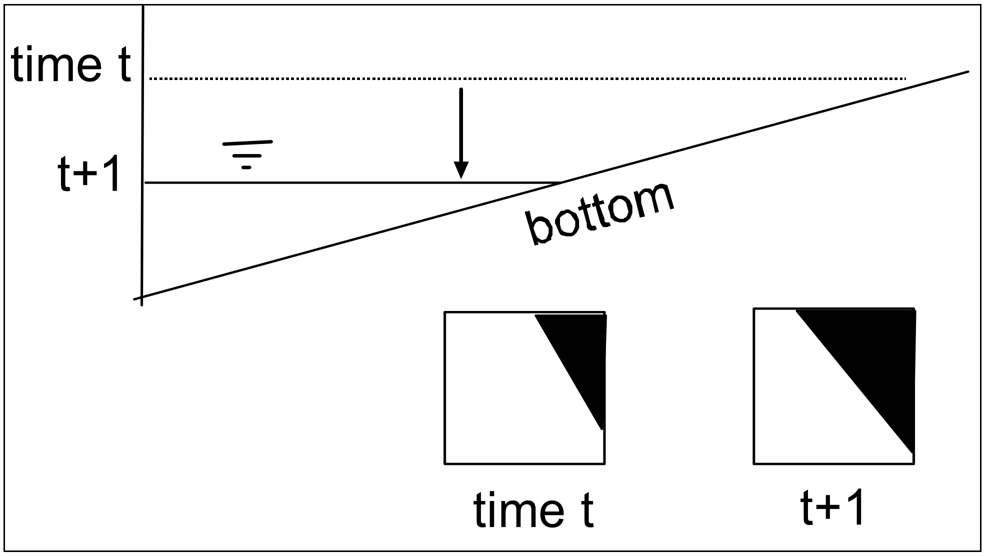

Figure 2.3: The Volume of Fluid (and Therefore the Boundary) Moving with the Water Level in a Computational Cell During Wetting and Drying.

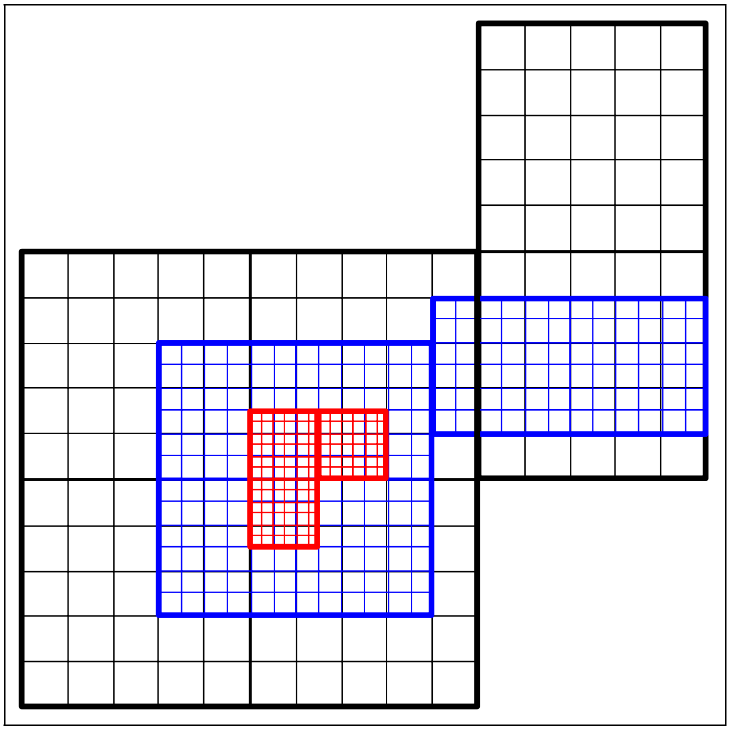

Figure 2.4: Adaptive mesh refinement.

2.3.2 Solution Techniques for Physical Processes

Aside from the grid, the numerical specifications for REALM provide for a great deal of flexibility. The design will allow the computational kernel (which describes the transfer of flux between cells) to be easily swapped, so that the solver to can be tailored to a particular flow regime or geographical region. Further swapping is also needed to investigate tradeoffs between model performance and accuracy in the context of parallel processing. Some of the initial solution techniques that will be incorporated in the prototype for particular physical processes are described below.

Advection

Advection refers to the transport of a property (momentum, concentration) by the local fluid velocity. In the transport equations for water quality, this is the main mechanism for carrying a constituent from one point to another. In the hydrodynamic equations, the cell flux may be generalized to include the wavelike propagation of information at the characteristic speed of the medium, which is faster than fluid flow (see the special discussion on *gravity waves* below).

Advection is notoriously difficult to model well, and the challenge to do so is a driving force in research on numerical methods. Only in the last 10 years or so have robust, efficient solvers existed for multi-dimensional advection. The prototype of REALM will use second-order Godonov methods with approximate Riemann solvers for calculating fluxes between cells such as the ones described in LeVeque (2002). These methods are compatible with 1D and 2D shallow water formulations, and are valid for solving any problem to which the shallow water equations are applied – from estuary problems to floods and levee breaks.

High-resolution Godonov methods are computationally demanding, although they scale well under parallel computation. If, under experimentation, the prototype methods prove too expensive, simpler flux calculations will be adapted using reputable schemes like the MacCormack methods for estimating intercell fluxes for hydrodynamics and QUICK or QUICKEST method for advection of passive scalars/constituents.

Gravity Waves

Special consideration must be given to the *gravity wave* term in the momentum equation, which describes the pressure force due to the water surface slope. In a shallow estuary, propagation of waves by this mechanism is much faster than the fluid velocity.

Numerically, the time step of the model must be fine enough so that gravity waves do not travel too far in one time step relative to the spatial mesh – a situation which causes instability in explicit models and inaccuracy in implicit models.[2] Some models (such as the local RMA and TRIM models) treat the gravity wave implicitly for stability and deliberately use a larger time step. This treatment is thought to be justified physically under gentle estuary conditions when water surface changes are small and gradual. It is known to speed up model performance when serial processing is used for the computations.

The economy and legitimacy of this simplification will be revisited in the REALM project, since the computational context (parallel processing) and variety of applications for REALM are being

---

[2] An extensive discussion of explicit vs. implicit numerical methods is beyond the scope of this document. As a generalization implicit methods tend to be computationally expensive per time step and stable over long time steps; explicit methods are computationally inexpensive per time step but require short time steps. The competing factors (computation per time step vs. length of time step) provide a performance vs. accuracy tradeoff that is context dependent.

developed are somewhat different than the circumstances that originally led to the implicit gravity wave treatment. In fact, Godonov methods calculate the flux into a computational cell in a manner that does not readily decompose into separate gravity wave and advection components.

Diffusion

Diffusion plays a number of roles in tidal problems. In the hydrodynamic equations, the transport of momentum by turbulence is often approximated using a diffusion analogy. In the water quality equations, diffusion is used to represent both eddy diffusion (mixing due to circulating currents) and longitudinal dispersion (apparent mixing in a one-dimensional channel due to variations in convection velocity over a channel cross-section).

Diffusion is usually modeled implicitly – diffusion has a mathematical character (parabolicity) that requires more stability than does advection. Since both advection and diffusion appear in the same equations, the operators are often split, with implicit methods such as the Crank-Nicolson or the $L_o - \text{stable}$ method of Twizell et al. (1996) used for the diffusion operator and explicit terms used for advection. REALM will initially use one of these two implicit methods for diffusion.

Source Terms

Source terms include a number of body forces defined at a point, such as friction, reactions of constituents, wind and other stresses, as well as sources and sinks of solute and flow. From a computational point of view these terms are mostly the same and only a few things matter, namely: 1) whether the sources are numerically "stiff", containing processes that change along very different time scales, or 2) whether they are singular (e.g., Mannings friction is not well-behaved when a water body is nearly dry).

Source terms can be split and treated separately, as can diffusion terms, and this will be the initial approach used in REALM. Higher-order Runge-Kutta explicit methods will be used for non-stiff source terms, particularly variants of Runge-Kutta that are known to be compatible with optimal control methods. The TR-BDF method, an implicit two-stage Runge-Kutta type scheme, will be used for stiff source terms.

2.4 Analysis and Management Tools

The crucial feature that will distinguish REALM from current models will be the incorporation of tools that will make decision support and policy analysis easier.

REALM will provide:

- ☐ Interactive "CD-player" user interaction. The user will be allowed to start, stop, rewind, adjust boundary conditions, and replay the time period.

- ☐ Model steering. The user may specify operating rules for boundary conditions and hydraulic devices that are managed adaptively (e.g. gates that are opened or closed depending on the state of the Delta such as water quality values).

- ☐ Automatic sensitivity calculations using the optimal control technique of adjoint modeling. Adjoint modeling is an efficient method of calculating the sensitivity of results to design decisions or to variations in boundary inflows (at every point in time).

- ☐ Data assimilation. The model will support statistical filters that can fold noisy field data into the model, achieving an appropriate balance between model and measurement error.

- ☐ Multi-objective analysis and visualization. Visualization tools will be provided not only to display the output of a single simulation on a map, but also to mine the data for trends and display tradeoffs between competing objectives, such as stage, releases, and exports.

- ☐ A data management model specifically designed for institutional modelers with GIS support. The data management will be based on an industry standard (ARC-HYDRO) database schema for storing hydrologic features, but will include innovations that make it easier to archive alternative versions of the same feature; e.g., a "historical" Stockton ship channel and a "dredged" Stockton ship channel.

Initially only a few of the features from this long list will be available to the user; however, the REALM design recognizes that the mathematical underpinnings of many of the features listed above are very similar. For instance, real time data assimilation models (Kalman filters) use the same "tangent linear model" as adjoint gradient calculations do. If these core abstractions are included in the design, a framework will result that can be filled out gradually and incrementally – without the daunting retrofit that would be required to add features later.

REALM will be interactive as the user will be provided an interface with controls much like the controls of a CD player. The user will have the ability to save a model state for a particular timestamp, run the model forward, rewind and try again with different boundary conditions. The interface will be visual, allowing the user to monitor the state of the model graphically. Another "interactive" feature is the support for user-designed operating rules. An example of this would be to ask the model to reduce pumping and/or operate gates when stage in the South Delta gets too low.

REALM will also make optimization and sensitivity analysis easier by incorporating adjoint model calculations into its design. Adjoint models can calculate in one pass sensitivities that would take vast numbers of "perturbation runs" to calculate by trial-and-error. REALM will not include hardwired solvers for particular optimization problems, but rather mathematical functionality that will make solving any large scale optimization problem easier. This functionality is exceedingly difficult to retrofit.

Finally, there is a class of decision support that lies partway between optimization and interactive modeling. This occurs under the discipline of multi-objective programming, where tradeoffs between hard-to-compare objectives are explored graphically. The user discovers solutions that are optimal or near-optimal according to some set of preferences, and then explores the tradeoffs between the objectives. The multi-objective algorithm provides guidance to avoid strategies that are inferior according to all of the objectives. These combinations are hard to avoid based on trial and error alone.

All of these interactive tools and a great many simpler applications require convenient, routine methods of initializing multiple simulations and collating or aggregating their results. The user might want to plot statistics over time from Monte Carlo experiments or display on an interactive map the differences between a base and alternative run. REALM will use Geographical Information System (GIS) software as its map interface for two reasons. First, it is well-developed, standard software that will greatly reduce development time. Secondly, and equally important, it will allow REALM to interact with other data and models much easier, especially biological and habitat data.

2.5 Collaboration

The REALM project ties together ideas from a number of different fields: water resources, numerical modeling, computer science, statistics, optimal control, and data management. DWR has expertise in a number of these areas, particularly in hydrodynamics, optimization, and aquatic systems. Nevertheless, a state-of-the-art project cannot be accomplished without significant contributions from experts in other fields. Areas in which outside assistance will be relied upon include:

- Infrastructure libraries for grid management and parallel computation.
- GIS for data management, grid development and visualization of results.
- Manipulation and visualization of time series data.

The idea of "infrastructure" for numerical models is not new, for there exists a veritable sub-industry of high performance infrastructure for numerical models, much of it locally produced by Federal Department of Energy labs and UC campuses. These libraries provide grid management, mesh refinement, embedded boundaries, and a great deal of solver support under parallelized solution schemes (domain decomposition and adaptive mesh refinement). For those versed in the underlying subject matter, the libraries have a parsimonious, elegant design, and provide hooks into other high quality open source resources. REALM will not only rely on these software infrastructure libraries; but also extend this infrastructure for the support of multidisciplinary water resources models.

In addition to providing data management for the model, GIS will also be the front end for the application. GIS contractors will provide the following resources for the model:

- Tools for the assembly and control of model input, including the definition of model components, assignment of boundary conditions, and interaction with the model.

- Grid development tools, including:
  - Algorithms and tools for channel cross-section development in 1D channels, and
  - Tools for creating and managing hierarchies of Cartesian grids with embedded boundaries (defined below) – matching the data structures used in LBL libraries.

- Archival storage closely mirroring the ARC-HYDRO standard for (spatial) hydrologic data. The storage system will use the layering system developed for DSM2 (see Ateljevich and Pranger, 2002), so that multiple versions of model components can be stored and combined to form alternative scenarios.

- A system for producing mapped displays of the model state, both to display results and for use in the model controller. The model state will be retrieved from a native format, probably based on HDF5, an open source format and input/output library that is very common in high performance computing applications.

- A system for selecting time series data from a map for display using time series visualization tools.

Actual visualization of time series will be handled in a separate project called VTools, which is a rewriting and extension of the DWR's Vista suite (Sandhu, 1998). VTools will also bind the archival capabilities of Informix and the Informix TimeSeries Datablade (currently being used by DWR's Division of Environmental Services, or DES) with the loading format used the model, which will be based on HDF5 or the open source data storage tool BerkeleyDB. HEC-DSS storage will not be used for REALM in any significant way, although conversion utilities will be provided for HEC-DSS files.

2.6 Work Flow

REALM is expected to be developed in two major phases.

Phase 1 of the project will emphasize the initial design and creation of a functional prototype model for flow and transport. The phase is scheduled to take about 15 months from the start of contracted development. During this time, LBL will adapt their grid infrastructure and numerical solvers for estuarine hydrodynamics and water quality mass transport. At the same time, GIS contractors will work to develop a consolidated Delta bathymetry and a 1D/2D grid development tool. Towards the end of Phase 1, testing will take place on the link between the LBL grid and the GIS storage system. The conclusion of Phase 1 will see a working version of REALM with a few of the easier features implemented, such as model steering.

In Phase 2, model performance will be addressed and REALM will evolve into a production model. An initial Bay-Delta calibration will be undertaken and particle tracking will be added. Early in Phase 2, REALM will be ready for calibration, which will be followed by the first production release of the model. In parallel efforts, interactive control and optimal control will be added during Phase 2, and work will begin on numerical solvers for 3D vertically stratified flow. A second release of REALM is expected to take place about 12-18 months after the first containing some or all of these features. Thereafter, development will proceed in small design cycles rather than large pre-defined phases.

2.7 Example Problems

A few example problems gleaned from interviewing managers and workers in the Delta should serve to illustrate the type of problems that REALM will solve.

2.7.1 Through Delta Facility

The Through Delta Facility (TDF) is a proposed inter-Delta diversion to supplement the Delta Cross Channel (Figure 2.5). It would divert water from the Sacramento River (near Hood), to the Central Delta (near the Mokelumne River system), with the intent of improving water quality in the Delta.

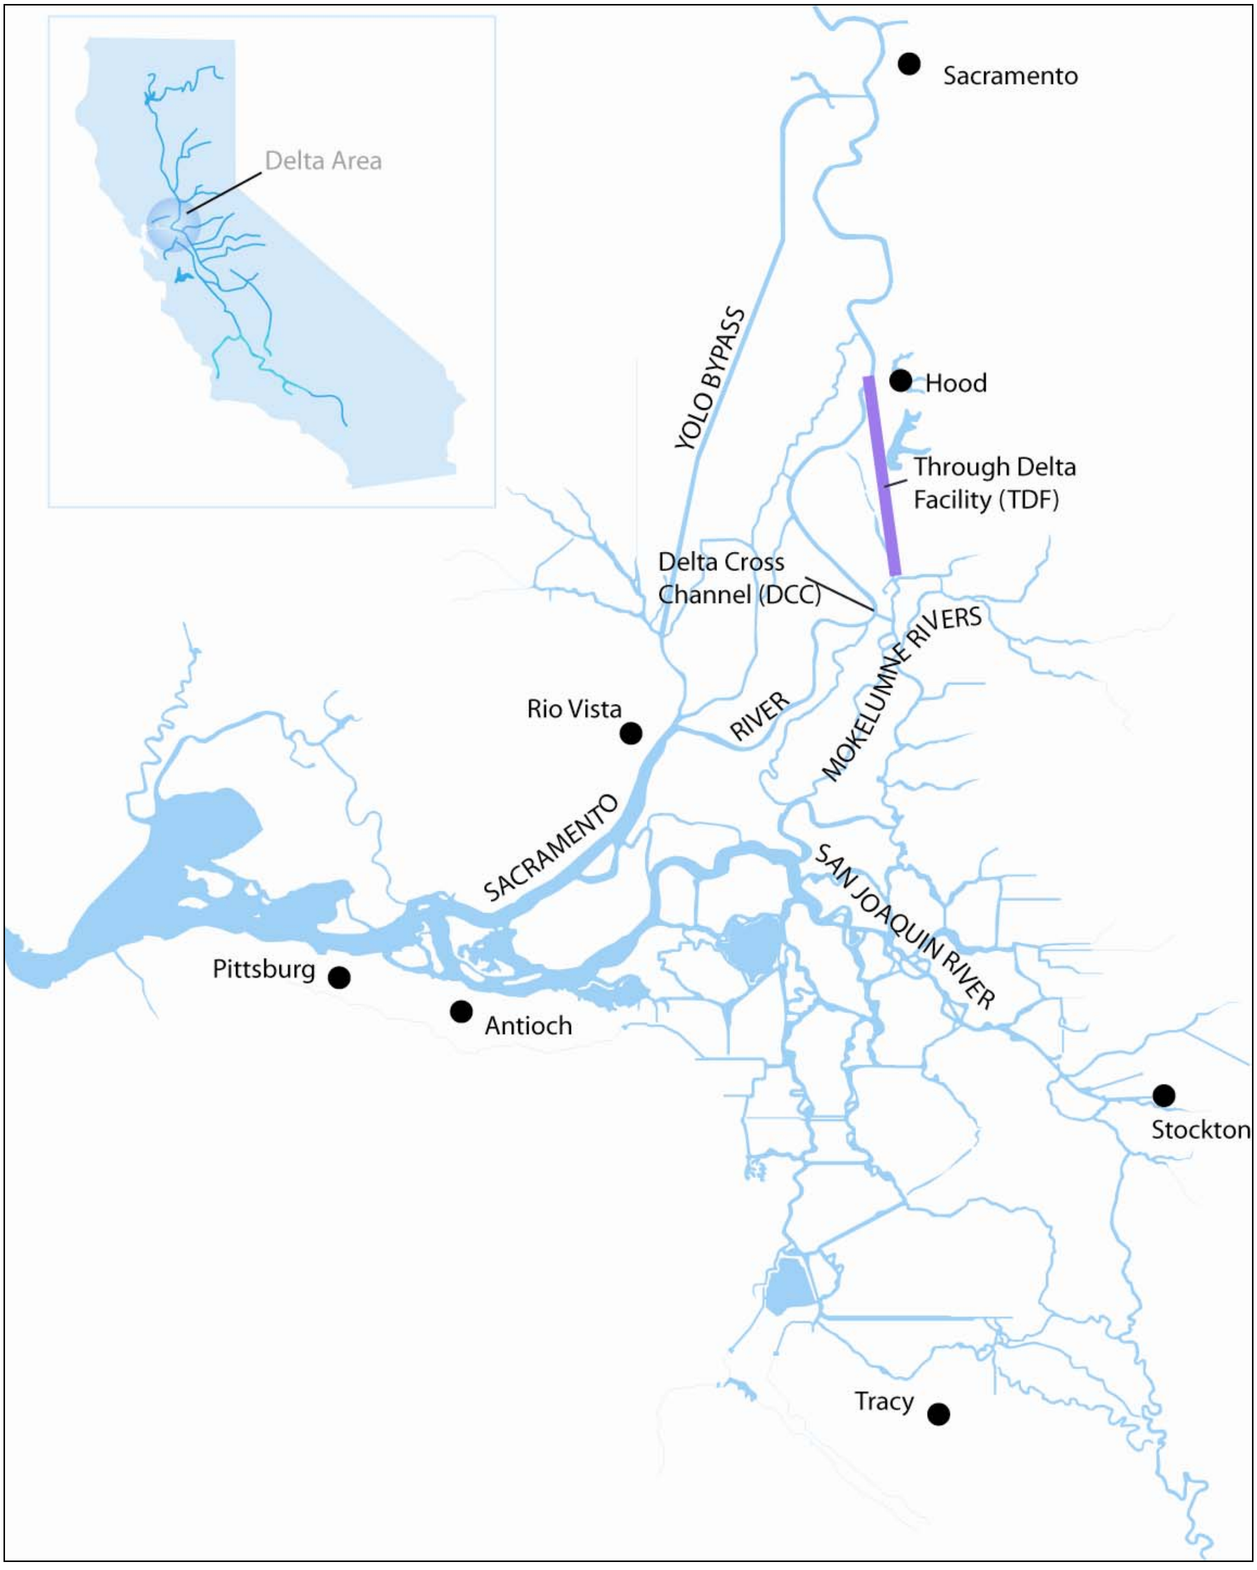

Figure 2.5: Through Delta Facility.

The Delta Modeling Section has conducted a series of studies using DSM2, a standard one-dimensional model, investigating different pumping rates and seasonal patterns of a TDF. The section was asked to run a study which would turn on the TDF when salinity at Rock Slough reached a certain level, then turn off the TDF when salinities dropped below another, lower target level.

Although this type of feedback from water quality to gate operation seems mechanically simple, on-the-fly manipulation of this sort is not available with DSM2 or any other model of the Delta now in common use. REALM's *model steering* feature will make it easy to make such a request by stipulating operating rules to guide the opening and closing of the TDF. Figure 2.6 shows how rules might be used to enforce salinity guidelines. The model anticipates the violation of user-imposed bounds and adjusts the TDF gates, export pumps, or upstream releases accordingly.

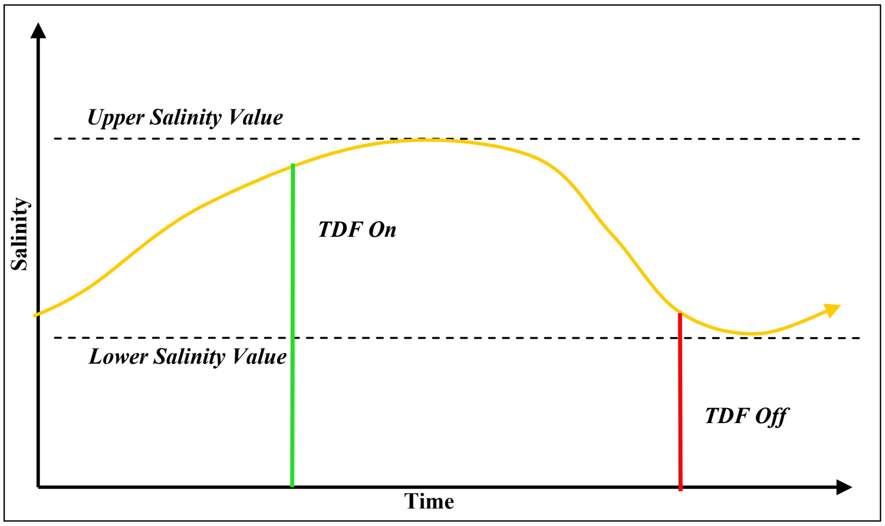

Figure 2.6: TDF-Controlled Salinity Within Range.

In the control scenario just described, the TDF or other device is operated to rules. The desired result is achieved, but perhaps at a price: too rapid a change in flow, too much water released upstream, or an unnecessarily large reduction in exports. If the controls (upstream releases and TDF flows) are considered to be continuously adjustable, REALM, with its *adjoint-supplied gradients*, could achieve the desired operation with an optimal water cost (Figure 2.7).

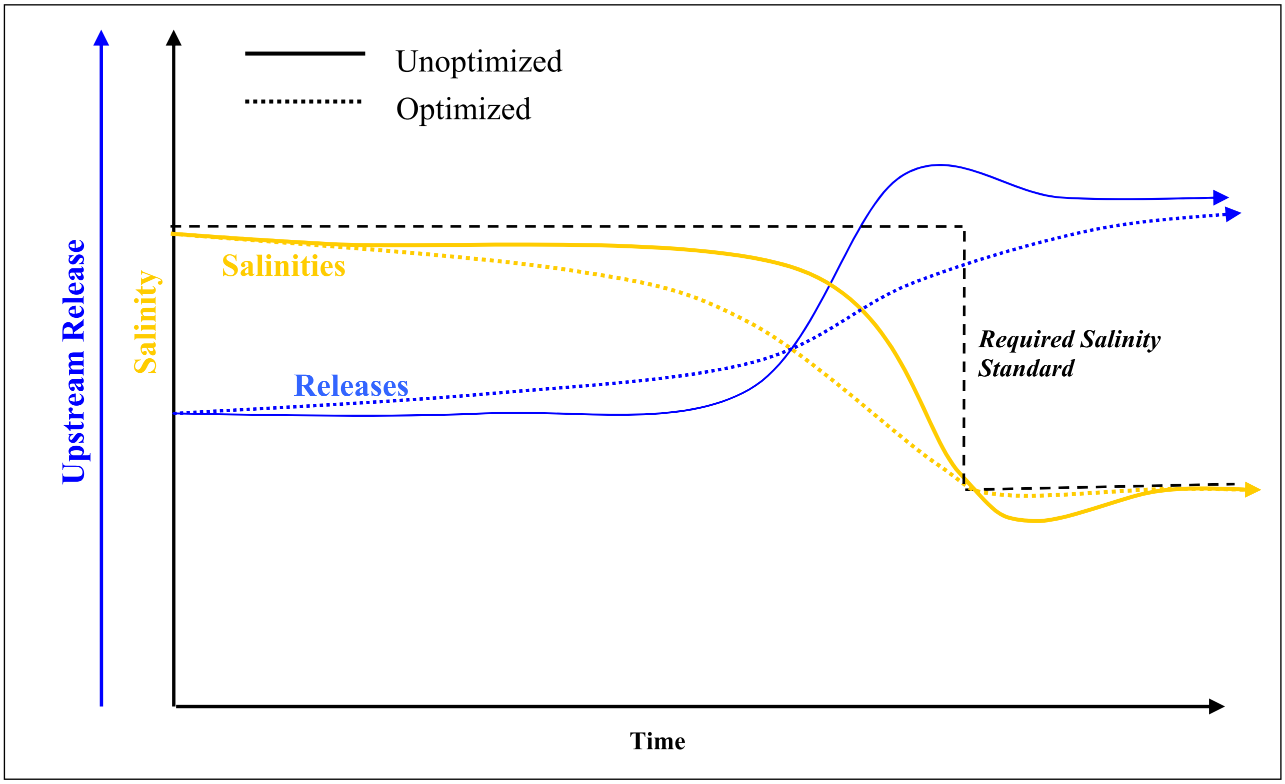

Figure 2.7: Unoptimized vs. Optimized Operation.

2.7.2 Particle Tracking

A number of Delta managers, engineers, and scientists were interviewed and asked to describe their model feature priorities. Virtually all of the environmental scientists interviewed said they needed accurate particle modeling and tracking. To do this, accurate velocity modeling in the junctions of channels is especially important, as well as accurate modeling of bays. This accuracy can be achieved only by using a 2- or 3-D treatment of junctions and bays, along with a fine-mesh grid in those areas.

Precise particle tracking will inevitably increase running time of the model. REALM's ability to vary the computational density of grid points and combine 1-, 2- and 3-D treatments, will allow the user to achieve the necessary level of accuracy where needed, without performing unnecessary computations within channels where high levels of accuracy would not be used.

2.7.3 Real-Time Adaptive Delta Management

Real-time, adaptive management of the Delta will play a more important role in coming years. Better and cheaper real-time data is available for traditional hydrodynamic and water quality values and biological and aquatic parameters, such as locations of endangered fish and counts. These data may enable adaptive management, but only if modeling tools are up to the task.

REALM will play a key role in this area because of its capability of real-time data assimilation. Data assimilation is the ability to incorporate observed data in a statistically correct manner into a model run. The observed data is compared to the computed data at the same time and location; inevitably the two values differ. Neither the observed or computed values are automatically accepted as “correct”. Instead, using a process called Extended Kalman Filtering, their respective errors are compared from physical and mathematical principles, and the model adjusted accordingly. Data assimilation does more than adjust the single model value at the location of the observed data; it also correctly modifies neighboring points in the model. In this manner the statistically best possible forecast is developed using all possible data.

2.8 References

Ateljevich, E. and T. Pranger. (2002). “Chapter 13: DSM2 Input Database and Data Management System.” *Methodology for Flow and Salinity Estimates in the Sacramento-San Joaquin Delta and Suisun Marsh. 23rd Annual Progress Report to the State Water Resources Control Board.* California Department of Water Resources. Sacramento, CA.

LeVeque, R. J. (2002). *Finite Volume Methods for Hyperbolic Problems*. Cambridge University Press.

Sandhu, N. (1998). “Chapter 8: VISTA (Visualization Tool and Analyzer).” *Methodology for Flow and Salinity Estimates in the Sacramento-San Joaquin Delta and Suisun Marsh. 19th Annual Progress Report to the State Water Resources Control Board.* California Department of Water Resources. Sacramento, CA.

Twizell, E.H., A. B. Gumel, and Arigu M.A. (1996). *Second order* L_0-*stable methods for the heat equation with time dependent boundary conditions. Adv. Comput. Math.,* ***6*** *(333).*