

# Integrated Computing Environments for Watershed Management

Layne T. Watson<sup>1</sup>; Vinod K. Lohani<sup>2</sup>; David F. Kibler<sup>3</sup>; Randel L. Dymond<sup>4</sup>; Naren Ramakrishnan<sup>5</sup>; and Clifford A. Shaffer<sup>6</sup>

**Abstract:** This paper describes a new computing paradigm known as a problem solving environment (PSE), explores the role this new paradigm might play in watershed management and land use change analysis, and compares PSEs to similar technologies such as decision support systems and geographic information systems. A prototypical watershed management PSE organizes and unifies the diverse collection of software typically associated with ecosystem models (hydrological, economic, and biological). A PSE provides a Web-based interface for potential watershed managers and other users to explore meaningful alternative land development and management scenarios and view their hydrological, ecological, and economic impacts. A brief description of a land use change analysis PSE being developed (called L2W, landscapes to waterscapes) is presented.

**DOI:** 10.1061/(ASCE)0887-3801(2002)16:4(259)

**CE Database keywords:** Algorithms; Watershed management; Land usage; Computation.

## Introduction

Watershed management models and systems come in many shapes, sizes, types, and names. All have limitations based on the system design and/or data. Many older research studies would focus on one particular aspect of a watershed, like hydrology, habitat quality for stream biota, and economic evaluation, of varying residential patterns. More recent research has focused on integrating several discipline-specific models into what has been called a *decision support system* (DSS). In many cases such systems incorporate a *geographic information system* (GIS) specifically targeted at watershed management. These two generic types of systems share a number of common features. In fact, some DSSs include a GIS component and, vice versa, some watershed GISs include decision support components. The perspectives and experiences of the research team members determine the features and terminology of the system.

One active area of research in computer science involves working with domain specific problems to create integrated prob-

lem solving environments (PSEs), which this paper attempts to carefully delineate from decision support and geographic information systems, as those terms are traditionally construed. PSEs were originally introduced in domains such as partial differential equations (PDEs) (Rice and Boisvert 1985; Houstis et al. 1998) and linear algebra (Gannon et al. 1998) where they provided high-level programmatic interfaces to widely used software libraries (Gallopoulos et al. 1994; Rice and Boisvert 1996). With rapid advances in high performance computing, GIS, software interfaces, computational intelligence, and networking, interest in PSEs has virtually exploded. Diverse applications in wood-based composite design (Goel et al. 1999), aircraft design (Goel et al. 2001), gas turbine dynamics (Drashansky et al. 1999), and wireless communications (Fortune et al. 1995) are now being viewed in the PSE framework. While these projects concentrate on developing domain-specific PSEs, considerable attention has also been devoted to developing generic tools for building PSEs. The software engineering of customizable architectures, leveraging the Web, supporting distributed, collaborative problem solving, and providing *middleware* constitute some of the enabling technologies. This paper describes specific watershed management system design issues from a PSE viewpoint, and discusses related research in a multidisciplinary context, including hydrology, economics, biology, geography, and computer science.

The purpose of this paper is to describe the critical features of a PSE for land use change analysis. While there is no doubt that the need exists for better models for all aspects of watershed assessment, including hydrology (flooding and erosion effects), biology (effects of contaminants and population changes); and economics (valuations resulting from land use changes and surrounding environment, economic effects on governments), the synergy resulting from integrating disparate models in a PSE will help leverage them in ways that best benefit planners and other observers. The whole (PSE) is greater than the sum of the parts (data, models, DSS, GIS), and the power lies in automating and leveraging human analysis and planning capabilities. A number of distinct aspects that should be part of a full-fledged PSE for wa-

<sup>1</sup>Professor, Depts. of Computer Science and Mathematics, Virginia Polytechnic Institute and State Univ., Blacksburg, VA 24061.

<sup>2</sup>Assistant Professor, Dept. of Engineering Fundamentals, Virginia Polytechnic Institute and State Univ., Blacksburg, VA 24061.

<sup>3</sup>Professor, Dept. of Civil and Environmental Engineering, Virginia Polytechnic Institute and State Univ., Blacksburg, VA 24061.

<sup>4</sup>Associate Professor, Dept. of Civil and Environmental Engineering, Virginia Polytechnic Institute and State Univ., Blacksburg, VA 24061.

<sup>5</sup>Assistant Professor, Dept. of Computer Science, Virginia Polytechnic Institute and State Univ., Blacksburg, VA 24061.

<sup>6</sup>Associate Professor, Dept. of Computer Science, Virginia Polytechnic Institute and State Univ., Blacksburg, VA 24061.

Note. Discussion open until March 1, 2003. Separate discussions must be submitted for individual papers. To extend the closing date by one month, a written request must be filed with the ASCE Managing Editor. The manuscript for this paper was submitted for review and possible publication on June 14, 2001; approved on February 8, 2002. This paper is part of the *Journal of Computing in Civil Engineering*, Vol. 16, No. 4, October 1, 2002. ©ASCE, ISSN 0887-3801/2002/4-259-268/\$8.00+\$0.50 per page.

tershed management are identified here, along with a rationale for the desirability of each point.

The design characteristics of a PSE are listed and discussed in the “Meaning of Problem Solving Environment” section. The “PSEs, DSSs, and GISs” section addresses the advantages and disadvantages of decision support systems and geographic information systems, contrasting them with PSEs. Each scientific application domain elicits unique considerations for the design and implementation of a PSE for that domain. The “Key Issues in Watershed Management” section addresses issues specific to watershed management and land use change analysis, and some previous DSS development efforts. Land use change analysis necessarily requires hydrologic simulation. The focus of the paper is still land use change analysis, but that entails a significant component (“Previous DSS/PSE Efforts...” section) on hydrologic tools. It also happens to be the case that many extant land use change systems focus on hydrology. The “Landscapes-to-waterscapes...” section offers a vision of the use of PSEs in watershed management in the future, and an example in the form of a brief description of a prototype PSE [landscapes to waterscapes (L2W)] for land use change analysis. The final section contains a technical challenge.

## Meaning of Problem Solving Environment

The purpose of this section is to offer a definition of a “problem solving environment” by listing and defining its attributes in the section directly following, and discussing the implications of those attributes for watershed management in a later section. As a starting point, take the following as a working definition (Rice and Boisvert 1996): “A problem solving environment is a computational system that provides a complete and convenient set of high level tools for solving problems from a specific domain. The PSE allows users to define and modify problems, choose solution strategies, interact with and manage appropriate hardware and software resources, visualize and analyze results, and record and coordinate extended problem solving tasks. A user communicates with a PSE in the language of the problem, not in the language of a particular operating system, programming language, or network protocol.”

### PSE Attributes

An object is defined both by what it *is* and what it is *not*. This section defines the key attributes of a PSE (what it is), and the “PSEs, DSSs, and GISs” section later elucidates differences between a DSS, GIS, and PSE (what it is not).

### Internet Access to Legacy Codes and Remote Compute Cycles

Legacy codes are large, existing, validated, and trusted computer programs that cannot reasonably be rewritten to interface cleanly with or run in a new computer environment. The term “remote compute cycles” refers to executing a simulation or graphics computation on a remote computer (ideally transparent to the user) rather than on the user’s local computer. Internet access means that all computation and interaction is done via a World Wide Web connection, rather than via a login or socket connection to a remote machine. Rather than create new simulations, PSEs can be more effective in modeling the effects of landuse change by integrating existing packages and software. Linking various models together via the Internet avoids platform depen-

ency issues, as users are not required to install the modeling codes on a compatible platform. Perhaps more important, by using a network-based approach, it is not even necessary that all of the models reside/run on the same platform, and PSEs can be envisioned as providing network-based “software services.”

### Interactive Visualization and Computational Steering

Interactive visualization constitutes viewing results of simulations graphically during or after the simulation run, with the user having some control over the content and format of the graphics during their production. Computational steering refers to *modifying the simulation itself* during the course of the simulation run, perhaps in response to some intermediate graphical output. Users of a PSE typically wish to visualize the output, rather than process raw data output by the models. Such visualization processes should be integrated seamlessly with the computational pipeline by the PSE. An important aspect of such integration relates to inlined simulation and visualization tasks (contained at the same level within a larger context). It can be argued that if one can identify specific processes (and/or subdomains) that are interesting, then computational resources could be steered toward these processes, while supporting other simulation tasks only insofar as to maintain the fidelity of the interesting phenomena. This concept of computational steering plays an important role in reducing the overhead associated with large-scale simulations.

### Scenario and Experiment Management

A scenario is a sequence of events such as simulation runs interspersed with data input and result storage decisions. Experiment management, which can significantly enhance scientific productivity, refers to archiving, retrieving, comparing, and mining the input and output from computer simulation runs. Recording scenarios in a database at a cognitive level meaningful to the user (above the raw model input data) can support experiment management, parameter tuning, and automated optimization.

### Multidisciplinary Support and Usage Documentation

Multidisciplinary support means explicit support for computer codes from multiple disciplines (e.g., hydrology, hydraulics, contaminant transport, economics, biology, urban planning), understanding that any user will not be expert in all of the disciplines. Usage documentation means providing expert advice, codifying the practical experience of disciplinary experts, for using all components of the system. Since the collection of models comprising a watershed assessment system are likely to be multidisciplinary in nature, a PSE must provide support to users who will not be expert in every (or any) aspect of the domain. One solution is to provide alternate interfaces to different aspects of the modeling subsystems to always be at the proper cognitive level for the user.

### Recommender Systems

A recommender is a computer program that functions as an automated reasoning assistant, guiding user decisions based on a database of facts, previous inferences, and user preferences. A full-fledged PSE will likely provide a rich collection of simulations for modeling various aspects of the problem. Unfortunately, the multitude of choices available can bewilder novice users. Recommender systems for PSEs serve as intelligent front-ends and guide the user from a high level description of the problem through every stage of the solution process, providing recommendations at each step.

### Collaboration Support

Collaboration refers to real-time interaction between geographically separated users, concurrently using the same PSE to work on the same problem. Decision makers often would like to either communicate their rationale to others, or work collaboratively with others during the planning process. While the ability to save and restore prior results can be used to provide asynchronous collaboration, ideally a PSE would allow multiple users at multiple sites to work together collaboratively and interactively. For instance, one user can create a scenario and display the results to others who can perform further analyses. Alternatively, two or more users can jointly set up scenarios. Together with component-based architectures, collaborative systems help realize the paradigm of "programming-in-the-large," where powerful programming abstractions harness widespread computing resources in an intuitive and transparent manner.

### Optimization

The precise mathematical meaning of optimization is intended here, whereby an objective function is minimized subject to constraints. Selecting a "best" configuration to balance competing goals within a watershed can be cast as a multiobjective optimization problem. A given run of a model is mathematically a function evaluation at a point in a multidimensional space. In essence, the goal is to supply to the model that vector of parameters that yields the best result under some figure of merit (or multiple figures of merit, for multicriteria optimization). As such, decision-making processes can often be improved by applying automated optimization techniques, rather than manually running a large number of simulations to explore the parameter space. Automated optimization techniques are quite sophisticated today, and are woefully underutilized by decision support systems in many disciplines, including watershed management. There exists great opportunity for significant improvement in the value of planning tools with relatively little development effort, since the state of the art in optimization tools far exceeds their current level of use in this domain. As mentioned above, many model users are currently spending large amounts of both human and computer time trying to do what amounts to optimization by hand. That time and computing resource would be better spent with the human acting at a higher cognitive level by describing the evaluation criteria, and then using automated optimization to seek out acceptable solutions that meet those criteria.

### High Performance Computing

High performance computing means the use of supercomputers (vector or parallel machines, clusters of workstations) when dictated by the problem. Many of the models used in watershed assessment require significant computing resources, such as a parallel supercomputer or an "information grid." PSEs can incorporate a computing resource management subsystem such as Globus (Foster and Kesselman 1997) or Legion (Grimshaw et al. 1994), and hide the details of accessing the necessary computational resources from the user. The transparent provision of the necessary computer resources to run a simulation in a reasonable amount of time (or inform the user of the intractability of the requested simulation) is an enormous boon to productivity; the time wasted moving code and data to different computers, and recovering from failed simulation runs due to insufficient resources, cannot be overestimated.

### Preservation of Expert Knowledge

Legacy codes (some 25 years old) continue to be used because they are trusted, they embody the practical experience and expert knowledge of their developers, and their limitations are well understood through decades of use. Preservation of expert knowledge refers to codifying knowledge in a computer program (PSE) in a form such that it can continue to benefit future generations. Like books in libraries, programs codify and preserve expert knowledge about the application domain. By using and preserving legacy codes, the expert knowledge embodied in the legacy codes is (indirectly) employed by the PSE. Yet, state-of-the-art codes in their native form are nearly impossible for nonexperts to use productively. By providing improved interfaces, tutorial help, usage documentation, and expert advice, either from knowledge culled from experts or by automatic inference and mining, PSEs can make legacy codes and knowledge more usable by nonexperts.

### Designed Extensibility

Designed extensibility means having modifiability and extensibility as an explicit design criterion, rather than as an afterthought. It is one thing to say that a DSS, by brute force, can be extended to support some capability. It is quite another to say that a DSS, by an explicit design intent, can be extended in some direction. With PSEs, as with any complicated computer system, scalability is always an issue. Adding more disciplinary analysis modules, or adding more capability to any aspect of the PSE (e.g., more visualization routines), must be accommodated as part of the basic design of the PSE, rather than as an afterthought. Further complicating scalability is the desire (perhaps at some future time) to use supercomputing and distributed computing, both of which place severe architectural constraints on the PSE design. The computer codes involved in a multidisciplinary PSE will be extremely heterogeneous, incompatible even, and the PSE architecture must gracefully deal with this reality.

### Pedagogical Uses

A pedagogical use of a PSE is its use as a teaching instrument. PSEs in domains such as watershed assessment can also help to improve education in all of the related disciplines. Students in environmental and civil engineering can more easily be made aware of biological and economic issues; likewise, biologists and economists can acquire sensitivity to technical issues in the other disciplines. Often the general public gets heavily involved in controversial zoning and planning decisions. Using PSEs, citizens could go online and learn about the various aspects involved in resource management decisions. They could evaluate for themselves the rationale for planning choices made in particular projects. Ultimately, a better understanding of the complex issues involved will benefit all parties.

### Implications of PSE Attributes

The power of a PSE derives from integrating existing packages and software. This is most conveniently done by providing *internet access to legacy codes and remote compute cycles*. This approach avoids platform dependency issues, as users are not required to install the system on a compatible platform, and it is not even necessary that all of the models reside/run on the same platform (Casanova and Dongarra 1997; Markus et al. 1997).

*Interactive visualization* is a sine qua non of any PSE. A significant extension of visualization is the concept of *computational steering* (Parker et al. 1997), whereby computational resources are steered during a simulation run toward specific interesting

processes (and/or subdomains), while supporting other simulation tasks only insofar as to maintain the fidelity of the interesting phenomena. Such steering is especially valuable for reducing the time required for large-scale watershed simulations.

PSEs should encourage users to experiment with various management options or scenarios. Such scenarios should be at a cognitive level relevant to the user; i.e., typically higher than the raw input demanded by the model. As each scenario is evaluated, the results can be recorded in a database for later retrieval, and for automated comparison to other scenarios. It is not uncommon for a typical user to run a model several times, with various combinations of input parameters, to generate output that meets some performance criteria. In some cases, users may conduct hundreds of experiments. Recording scenarios can thus aid in *scenario and experiment management* (Ioannidis et al. 1996), parameter tuning, and automated optimization. In addition, having the ability to perform *data mining* (Ramakrishnan and Grama 1999, 2001) with respect to desired characteristics provides powerful analysis capabilities for what-if scenarios.

The use of a PSE for large scale multidisciplinary research problems requires *multidisciplinary support and usage documentation*. This most likely will require alternate interfaces to different aspects of the modeling subsystems to reflect various levels of expertise, consistent with the notion that scenarios should be at the proper cognitive level for the user. Typically, expert users desire more detailed control of models, while novice users will wish to control only the coarse details, and need the maximum amount of guidance on reasonable parameter settings for models. The simulation interface could provide recommendations on reasonable interactions of parameters, or on which submodels to use in particular circumstances. Such advisory support regarding parameters is an integral aspect for the practical utility of PSEs. As an example, WBCSim (Goel et al. 1999) provides a detailed description, with typical values and ranges, for every parameter in the models it supports.

*Recommender systems* for PSEs (Ramakrishnan et al. 1998), which guide the user from a high level description of the problem through every stage of the solution process, providing recommendations at each step, are especially important for multidisciplinary watershed management systems, with diverse users. Since planning for land use and watershed management is typically a collaborative process, a PSE should provide real-time *collaboration support* to allow multiple users at multiple sites to work together collaboratively and interactively.

Decision-making processes can often be improved by applying automated optimization techniques, rather than manually running a large number of simulations to explore the parameter space. That time and computing resource would be better spent with the human acting at a higher cognitive level by describing the evaluation criteria, and then using automated optimization to seek out acceptable solutions that meet those criteria. Optimization and high fidelity simulations often require access to *high performance computing*, meaning a parallel supercomputer or an information grid. The transparent and automatic provision of the necessary computer resources to run a large scale watershed simulation in a reasonable amount of time (or inform the user of the intractability of the requested simulation) is crucial for productivity.

Like written documents, legacy watershed programs codify and preserve expert knowledge. By providing expert advice, either from knowledge culled from experts or by automatic inference and data mining, PSEs can make legacy codes and knowledge usable by the diverse clientele of a watershed management

system. Thus, PSEs play a social role in the *preservation of expert knowledge*.

With watershed management PSEs, scalability is a major design issue. Adding more disciplinary analysis modules or more capability to any aspect of the PSE (e.g., more visualization routines) must be accommodated in the basic design of the PSE. The ultimate necessary use of both supercomputing and distributed computing places severe architectural constraints on the PSE design. The computer codes involved in a multidisciplinary PSE will likely be heterogeneous and incompatible, so the PSE architecture must gracefully deal with this reality by having *designed extensibility*.

The usage documentation and expert disciplinary advice for novice users make PSEs also well suited for *pedagogical uses*. For public policy issues such as urban planning or land use change analysis, citizens could learn about the various aspects involved in resource management decisions by going on-line to a Web-based PSE. They could evaluate for themselves the assumptions, constraints, and rationale for planning choices made in particular projects.

## PSEs, DSSs, and GISs

It is instructive to understand the cognizant technical issues arising from the (respective) parent communities of GIS, DSS, and PSE research. GISs represent, analyze, manage, and integrate data that are expressed on a spatial scale. The study of GISs includes issues such as linking databases, cartographic tools, query processing, and visualization of results. DSSs, with roots in business analysis and decision making, emphasize the role of an integrated system in supporting a broader array of functions such as logistics, site design, planning, market studies, and demographic analysis. Many commercial DSSs support GIS functions for domains such as environmental decision making. The PSE thread of research originally arose from the need to provide high-level access to legacy scientific codes. It has now expanded to include automatic selection of solution components, parameter optimization, experiment management, collaborative composition, and visual problem-solving capabilities. While it is conceivable that one single system could fit all of these definitions, GISs emphasize the *nature of data and information* (spatial), DSSs emphasize a *functionality* (analysis, planning, and decision making), and PSEs emphasize a *problem domain* (such as watershed management). In summary

GIS = spatial data + set of programs to access, manipulate,  
and visualize spatial data

DSS = decision support + analysis

PSE = simulation codes + networked access  
+ high-level interface + intelligence

The words in these equations are content laden. While a GIS may involve a legacy simulation code, no GIS supports a large scale fluid dynamical PDE solution of river flow with bank erosion and sediment transport, requiring a supercomputer, which is the import here of “legacy code and remote compute cycles.” “Networked access” includes management of remote and distributed compute cycles, and “intelligence” includes recommender systems and optimization—all far beyond the scope of GISs and DSSs.

## Relation of PSE to DSS

Decision support systems have been defined (Sol 1983) as “applying information systems technology to increase the effectiveness of decision makers in situations where computers can support and enhance human judgment.” In certain domains (clinical decision making, for instance), DSSs are used synonymously with expert systems that support the problem-solving process. In the area of water resources development, Reitsma et al. (1996) define DSSs as “computer-based systems which integrate state information (i.e., data representing water resource system’s state at any point of time), dynamic or process information (i.e., principles governing resources behavior over time), and plan evaluation tools (i.e., utility software for transforming raw system data into information relevant for decision makers) into a single software implementation.” In other domains (Kant et al. 1992), the scope of DSS is extended to include problem formulation. By modeling a complete problem-solving process, the PSE paradigm provides powerful conceptual abstractions that go beyond traditional DSS facilities. Four increasingly complex levels of functionality can be identified.

1. *Management of the execution environment.* PSEs solve persistent software infrastructure issues underlying the integration of software from multiple disciplines and sharing of software among multiple stakeholders, collaboratively. Technologies such as the software bus (Purtilo 1994) and the Symphony framework (Shaffer et al. 2000) for manipulating remote resources allow seamless access to high-performance resources such as grids and clusters of workstations. For example, a multidisciplinary and geographically distributed team of economists, planners, hydrologists, and biologists could investigate the effect of new settlement patterns on the phenomena of interest in a concurrent fashion. Support for this feature in DSSs is nearly absent and usually hardwired for a particular system. For instance, users do not have the ability, at run time, to collaboratively involve a new stakeholder in the problem-solving process.
2. *Experiment management.* By recording problem-solving sessions as entities in a database, PSEs allow the query, composition, and management of hundreds of simulations in a high-level manner. They can compile high-level specifications such as [Design simulations for the Roanoke watershed using the updated fisheries model from Marcus, but retaining the residential settlements of 1998.] Experiment management is now facilitated by powerful query optimization algorithms that can selectively reorder operations to address such questions efficiently. Two DSS-like systems that support this feature are the Sequoia project (which is just an application layer over a spatial database), and the ZOO experiment management system. Their domains are far detached from watershed management.
3. *Reasoning about models.* PSEs recognize the multitude of choices for modeling and simulating various phenomena (multiscale, multidisciplinary). By organizing a database of simulation runs (see above), and using technologies such as data mining, they can identify the most appropriate choices for models, given performance constraints. This feature is absent in all of the DSS systems surveyed here; limited reasoning is usually provided as hand-coded expert system rules (which have to be tuned periodically), or well calibrated system parameters (from published literature).
4. *Abstractions for problem solving.* The final aspect is the provision of conceptual and visual problem-solving capabilities that integrate design with analysis, for instance. In the wa-

tershed context, PSEs can take descriptions of scenarios, compose experiments automatically, farm out simulations on remote machines, perform optimization with respect to selected criteria, and repeat this in a grand loop when there are multiple criteria.

Some DSSs capture a recurring scenario in watershed assessment, and provide add-on modules that mimic such functionality. Creating support for new scenarios is often impossible. In summary, a PSE could thus be used for decision support; however, its emphasis on integrating model formulation, algorithm selection, representation, reasoning, data analysis, and experiment management provides a more holistic and scenario-based approach to watershed assessment and land use analysis.

## Relation of PSE to GIS

A GIS is a tool for doing spatial analysis. With its beginnings in diverse fields including cartography and surveying, remote sensing, and urban planning, computerized mapping and analysis tools date back to the 1960s. The 1980s saw the first successful commercial tools that have evolved into current GIS systems. A GIS is normally recognized to have the following components to a greater or lesser degree (Burrough 1986): “raster and vector data structures, modules for data input, verification, storage, and output, digital terrain models, methods of spatial analysis and modeling, and methods of classification and interpolation.”

The emphasis in a GIS that sets it apart from a computerized atlas is its ability to perform specific spatial analysis tasks. The ability to perform such tasks quite naturally leads to the ability to act to some degree as a decision support system. Examples of analysis tasks that can be expected from a typical GIS include interpolation from point data; data input and verification tasks including registration aids for joining maps or images to maps; scaling and projection tools; tools for calculating distances, sizes, and volumes; best paths routing; intersection of polygons versus polygons, lines, and points; classification of objects by attributes; locating nearest neighboring objects; and analysis related to data quality and errors.

Given these capabilities, it is quite clear that a GIS will play an important role in any decision support system or problem-solving environment for watershed management. The roles of a GIS in such a system include

- *Input/output.* User input will often be in terms of locations on a map. Visualization of results may take place in the context of a map.
- *Modeling and analysis.* A watershed PSE requires many models to calculate the necessary outputs for decision making. Some, but not all, of these models can be implemented directly within a standard GIS.

A GIS itself may have sufficient abilities to classify as a PSE for restricted domains whose analysis is strictly of spatial data. Certainly in the context of watershed analysis, a GIS itself is not sufficient to classify as a PSE. In particular, without extending the definition of GIS to the point of uselessness, the GIS is certain to be lacking the following features: domain-specific models related to economics, biology, and hydrology; tools for optimization (though optimization tools might be incorporated into a standard GIS for various spatial analysis tasks), experiment management tools, and recommender and data mining tools.

Of course, in principle, all of the PSE features listed in the section “PSE Attributes” could be added to a GIS, but then it would not be a GIS. A central contention of this paper is that

distinctions, based on definitions and common usage, between DSS, GIS, and PSE are meaningful, significant, and useful.

## Key Issues in Watershed Management

The key issues in watershed management and land use change assessment vary widely, depending on the question to be addressed, and on the size and complexity of the watershed under investigation. At the very small municipal end, local officials are often concerned about controlling unregulated storm water and non-point-source generation in developing portions of their drainage system. Their immediate concerns often take the form of how to regulate subdivision construction in such a way as to minimize runoff and sediment releases. Or they may be concerned about meeting state and U.S. Environmental Protection Agency (USEPA) mandates for National Pollutant Discharge Elimination System permitting, which impose standards for monitoring and inventory of their storm sewer systems. Here a common issue is whether to build larger regional storm water management facilities or smaller on-site facilities, which retain the excess runoff and its associated pollutants at the site where they are produced. Should a municipality take responsibility for long-term maintenance of the larger facilities, and if so is a storm water utility tax the best means to meet maintenance expenses? What are the impacts of zoning restrictions on downstream receiving waters? Does cluster residential and so-called low-impact residential development produce a measurable benefit to aquatic systems? If so, what are the adverse effects on land valuations? How can conservation districts and preservation land trusts be incorporated into regional master plans for managing long-term growth, and what are their impacts on sensitive receiving waters?

In larger metropolitan areas, this concern manifests itself in the form of combined sewer systems and the control of combined sewer overflows to sensitive receiving streams. Here, a micro-scale level of diagnostic hydraulic simulation of individual pipe elements, interceptors, and underground storage devices is required as part of the redesign process. In larger nonurban watersheds, specific concerns take the form of correcting problems of impaired streams that have made the state list because the concentration of some pollutant, like fecal coliform, has exceeded some standard. Here, the key issue is often which total maximum daily load (TMDL) tool to select as a DSS and what scale of application is needed in order to delineate pollutant source areas and develop an equitable policy for allocating the total maximum pollutant load among the stakeholders in the watershed.

In the case of very large systems, like the 166,000 km<sup>2</sup> Chesapeake Bay watershed, the concern is for unregulated releases of sediments, phosphorus, and nitrogen, which have the potential to destroy the shellfish industry in the bay itself. Once again, a complex hydrologic simulation tool is brought to bear to identify feasible solutions. A similar approach is taken in the case of large watersheds experiencing rapid urban growth. The municipal/regional planning agency often wants to know what the impacts of full build-out in the watershed will mean in terms of increased flooding, decreased groundwater recharge, and lowered in-stream baseflow, all of which have major importance for the aquatic habitat and overall stream quality of the region. Will the use of streamside riparian buffers and stream restoration projects offset the adverse impacts of prior development? Or has the stream surpassed a natural threshold beyond which no remediation is possible?

Often the critical issue facing the decision maker is simply the quantity of water, such as in the case of flood control in the upper

Roanoke River system or in the James or Neuse River basins. The issue here is often for long-term control of future floods via engineered facilities or perhaps for the use of real-time flood forecasts and improved flood warning systems, leading to better emergency action plans. This issue has emerged with greater focus because of recent U.S. Congress actions through the Federal Emergency Management Agency and the National Flood Insurance Program (NFIP) to restrict development in flood-prone areas. Grant moneys are often available to communities for devising plans that will reduce their vulnerability to natural disasters such as flooding (e.g., Project Impact program). In times of low rainfall, the reverse of this problem shows up in the form of contingency drought planning for emergency water supply. Perhaps the issues of watershed management are addressed most completely at the river basin commission level. As an example, the Susquehanna River Basin Commission, formed in 1971 by federal-state entities in five states, assumes responsibility for floodplain management, water supply, water quality, watershed protection, recreation, and fish and wildlife. Clearly, the charge to the commission requires a comprehensive database and a decision support system to evaluate impacts and risk associated with changes in this 71,000 km<sup>2</sup> basin.

The water resources planning literature is replete with examples of DSS models, usually in the form of simulation models having diagnostic capabilities, combined with a GIS. The question at hand is how the use of a more advanced PSE tool, that goes beyond a DSS, can add something of value to the decision-making process. A brief review of past DSS/PSE efforts in watershed management follows.

## Previous DSS/PSE Efforts for Watershed Planning and Management

Several notable attempts have been made in recent years to develop DSSs and PSEs for watershed assessment and management issues. The brief discussion here represents the key studies and latest developments known to the writers. Davis et al. (1991) developed a DSS consisting of three modules (namely, a policy module, a catchment module, and a query module) to examine the effects of potential land use and land management policies on water quality in South Australia. The policy module allows the user to build up a suite of policies, and the catchment module estimates the effects of these policies on total phosphorus, total nitrogen, and turbidity levels in a catchment under consideration. The query module allows the user to see the results of the simulation. To support everyday management of the Tennessee Valley Authority (TVA) river, reservoir, and power resources, a DSS named TVA Environmental and River Resources Aid has been developed (CADSWES 1993). This system provides the various geographically distributed TVA departments with historic, (near) real time, and (estimated) future information on the status and trajectory of the TVA water and power systems (Reitsma et al. 1996). AquaTool, a DSS developed in Spain for water resources planning and operational management in complex basins, is composed of several modules that are linked through geographically referenced databases and knowledge bases. The modeling capability includes basin simulation and optimization modules, an aquifer flow modeling module, and two modules for risk assessment. This DSS has been used by two river basins agencies in Spain for the management of water resources (Andreu et al. 1996).

At the hydrologic simulation level, the BASINS model released by the USEPA in September of 1996 integrates a GIS, a

**Table 1.** Features of Some Web-enabled DSSs

Category	L-THIA	Digital meadowlands	Management oriented watershed simulation environment	WFAT	HyDSS	Online hazard map	PLM	L2W PSE
Scope	Land use, hydrology	Land use, zoning	Watershed, multidisciplinary	Floodplain management	Hydrology	Hazard mapping	Watershed, multidisciplinary	Watershed, multidisciplinary
Transferability	Flexible	Site specific	— <sup>a</sup>	Site specific	Flexible	Flexible	Site specific	Site specific
Development stage	Intermediate	Initial	Proposed	Final	Final	Final	— <sup>a</sup>	Initial
Web access	Unrestricted	Restricted	— <sup>a</sup>	Unrestricted	Restricted	Unrestricted	Unrestricted	Restricted
User interface	Novice/expert	— <sup>a</sup>	Novice	Novice	Expert	Novice	Expert	Novice/expert
Reporting mechanism	Graphical/tabular	Graphical/tabular	— <sup>a</sup>	— <sup>a</sup>	Tabular	Graphical	— <sup>a</sup>	Graphical/tabular
Partners/funding agencies	EPA	NASA	— <sup>a</sup>	NASA/Stennis Space Center	NASA	ESRI/FEMA	NSF/EPA	EPA
GIS support	Yes	Yes	Yes	Yes	— <sup>a</sup>	Yes	Yes	Yes

<sup>a</sup>Not available.

national watershed database for major basins, and comprehensive environmental modeling tools into a single package for performing watershed assessment and water quality analysis. The primary use of this integrated watershed modeling framework is non-point-source management and development of TMDLs (<http://www.epa.gov/OST/BASINS/>). A number of European organizations have developed a DSS called WaterWare to assist river basin commissions on decisions for the efficient management of water resources in terms of both quantity and quality. The DSS consists of a GIS and a database management system, coupled to a number of analytical components, including demand forecasting, water resources planning, groundwater pollution control, surface water pollution control, and hydrological processes. The DSS has been applied to two river basins (in England and Mexico) to address the problems of water resource assessment, reservoir site selection, decontamination of groundwater, estimation of sustainable irrigation abstractions, and derivation of required effluent quality standards (Fedra and Jamieson 1996).

Dunn et al. (1996) describe the hydrology component of the NERC-ESRC Land-Use Program (NELUP) DSS with the objective of predicting the impact of agricultural land use change at the river basin scale. The model components of NELUP represent agricultural economics, ecology, and hydrological regimes of the basin. The hydrology models within the DSS are capable of analyzing hydrological effects at the catchment scale and at the field scale. Because of the complexity of various land use change problems, the NELUP DSS cannot be used by nonspecialists without assistance. Another DSS (Leavesley et al. 1996) for water and power management is called the modular modeling system (MMS). The MMS uses a master library that contains compatible modules for simulating a variety of water, energy, and biogeochemical processes. The GIS interface of the MMS is developed to facilitate model development, parameterization, application, and analysis. Typical applications of the MMS include management of multireservoir river systems within the constraints of competing water users and selected environmental constraints, such as water temperature limits or fisheries habitat needs. Simonic and Bender (1996) discuss the concept of a collaborative planning support system (CPSS) in water resources planning. This concept involves integrating available computer technologies with modeling and analysis tools in a user-friendly environment, thereby enhancing communication between the proponent for resource development and affected parties. The CPSS does not

provide solutions, but empowers participants by identifying areas of common understanding, encouraging them to explore solutions and reach a consensus. Osmond et al. (1997) developed a DSS with the following objectives: to transfer information to watershed managers for making appropriate land management decisions, to assess non-point-source pollution in a watershed based on user supplied information and decisions, and to evaluate water quality effects of alternative land treatment scenarios. Chen et al. (1999) discussed the development of a DSS, called WARMF, that includes a watershed simulation model, a database, a consensus building module, and a total maximum daily loads module that allows the calculation of TMDL for various pollutants within a river basin. The methodology is demonstrated by application in the 13,000 km<sup>2</sup> Catawba basin extending from North Carolina to South Carolina. This work has been described more recently by Weintraub et al. (2001).

While all of the above DSSs use computer models and link appropriate simulation modules via GIS, none of these are Web accessible by design. Studies are under way at a number of places to examine the potential of Web-enabled watershed management tools. Voinov and Costanza (1999) discuss the potential of Web-enabled watershed management tools in delivering scientific findings and information to stakeholders and in linking stakeholders together, thereby providing for collective decision making. The authors attempted to illustrate the potential of the Web-enabled approach by application to the Patuxent River. They concluded that it is not the amount and quality of information that are crucial for the success of watershed management, but how well that information is disseminated, shared, and used by the stakeholders.

A critical comparison of some Web-enabled hydrology related systems is given in Table 1. Most of these systems appear to be at a preliminary stage of development, with restricted Web access (Table 2), and include GIS support for the visual display of output results. Pandey et al. (2000) discuss a Web-enabled computer modeling tool, long-term hydrologic impact assessment (L-THIA), that is designed to assist urban and regional planners with the assessment of how land use changes impact long-term hydrology and non-point-source pollution in a watershed. The hydrologic computations within L-THIA are highly simplified and cannot be used to determine impacts on storm peaks or effects during a specific year. Therefore, the L-THIA authors do not recommend the use of the model for storm water drainage design or other urban infrastructure planning concerns.

**Table 2.** Web URLs for Systems in Table 1

System name and reference	Web URL
Long-term hydrologic impact assessment (L-THIA), Pandey et al. (2000)	<a href="http://danpatch.ecn.purdue.edu/~sprawl/LTHIA">http://danpatch.ecn.purdue.edu/~sprawl/LTHIA</a>
Digital meadowlands, Artigas et al. (2001)	<a href="http://cimic.rutgers.edu/digitalmeadowlands">http://cimic.rutgers.edu/digitalmeadowlands</a>
Management oriented watershed simulation environment, Westervelt (2001)	— <sup>a</sup>
Web-based floodplain advisory tool (WFAT), Sugumaran et al. (2000)	<a href="http://www.cares.missouri.edu/dem">http://www.cares.missouri.edu/dem</a>
Hydrology decision support system (HyDSS), Ram et al. (2000)	<a href="http://krishna.bpa.arizona.edu/HyDSS">http://krishna.bpa.arizona.edu/HyDSS</a>
Online hazard map, ESRI/FEMA	<a href="http://www.esri.com/hazards">http://www.esri.com/hazards</a>
Patuxent landscape model (PLM)	<a href="http://kabir.cbl.umces.edu/PLM/Welcome.html">http://kabir.cbl.umces.edu/PLM/Welcome.html</a>
Landscapes to waterscapes: A problem solving environment (L2W PSE), Rubin et al. (2000)	<a href="http://landscapes.ce.vt.edu/">http://landscapes.ce.vt.edu/</a>

<sup>a</sup>Not available.

Ram et al. (2000) describe a prototype state-of-the-art Web-based hydrology decision support system (HYDSS) that provides a comprehensive environment for information integration and analysis. The HYDSS user interface presents data and models like vegetation, soil, rainfall, and runoff that are directly or indirectly linked to runoff simulation. It appears from the description of HYDSS that simplified event based runoff generation approaches were used for hydrologic modeling. LUCAS (Berry et al. 1996) is a PSE integrating market economics, transportation costs, population density, and ecological behavior with an advanced graphical user interface for land use scenarios and habitat changes. Intended for socioeconomic modeling, LUCAS uses Markov models derived from time series data and expert opinion, so its simulations are random events predictive only in a probabilistic sense. Rubin et al. (2000) discuss a prototype PSE called L2W that can be used for multidisciplinary analysis of watershed development actions. A brief description of L2W follows.

Other related work includes that of Carpenter et al. (1999), DeBarry and Quimpo (1999), Grimshaw et al. (1998), and Kittle et al. (1998).

### Landscapes-to-waterscapes PSE—Vision for Future

To illustrate further the PSE features defined earlier, this section describes briefly a prototype PSE, L2W. A detailed case study of L2W is described elsewhere (Rubin et al. 2000). The L2W PSE project is built on a foundation of modeling capabilities in hydrology, economics, and stream ecology, supported by GIS data input layers and a user interface for selection of model function and land use scenario generation. The hydrologic simulation is done by the computer program HSPF (Bicknell et al. 1997), a comprehensive model of important surface runoff and subsurface flow processes. The economic land valuation model is used to evaluate shifts in municipal tax revenue and fiscal costs associated with land development. The ecologic assessment model of fish health is invoked to display the mean metric score map associated with a given land use scenario first simulated by HSPF.

A prototype application of the L2W PSE has been in the 145 km<sup>2</sup> Back Creek subwatershed of the upper Roanoke River basin in southwest Virginia. For the purpose of hydrologic simulations using HSPF, the subwatershed was divided into 10 land segments, with segments 1 and 2 representing the headwaters and segment 10 representing the mouth of the subwatershed. To identify the

developable land within the subwatershed, GIS layers representing already developed or protected lands, land with a slope exceeding 20%, water bodies, or floodplains were overlapped. About 50% of the subwatershed area was found to be available for development. The PSE provides access to the underlying models through a high level task oriented interface. The user chooses particular types of settlements (high density, medium density, medium density cluster, or low density) and places these settlements in particular locations within the watershed to develop a settlement scenario. A user creates a what-if scenario by choosing the new settlement pattern using the development scenario buttons, after which the predevelopment and postdevelopment land use distributions are calculated and provided to the user. This information is then used by the calibrated HSPF model that simulates the effects of development on various hydrologic parameters, including total annual runoff, selected storm peaks, and groundwater recharge. Detailed graphical output from the HSPF simulation can be shown at the land segment and subwatershed levels; Rubin et al. (2000) contains examples and an analysis of these hydrologic graphs. Hydro-meteorologic data for water years 1995–1998 are used to evaluate the hydrologic effects of different settlement scenarios. The impact of the development on hydrologic variables such as total runoff volume and groundwater recharge can be shown at the outlet of the subwatershed and at the outlet of the land segment. Similarly, the PSE has the potential to evaluate the effects of development on land values, public expenditures, and tax revenues. The L2W economics model specifies a complex regression equation, which requires a great deal of input from the GIS and its associated spatial data layers. In addition, the L2W PSE will have a fisheries model that relates the land use in zones of influence to the health of fish populations in those areas. These models have been developed using field data collected in and around the study site.

Hydrologic output is temporal, and is used to assess economic and ecological effects over time. Space limitations preclude a discussion of these temporal effects, which would be part of a detailed L2W case study—a future paper.

The elements described above provide a perspective on the contents of a full-blown PSE. In short, there must be GIS data layer capability; there must be a hydrologic simulation at spatial and temporal scales sufficient to fully represent the potential impacts of land change; there should be economic evaluation of costs and benefits incurred; finally, there should be some in-stream assessment of impacts based on the likely fish response.



The designed extensibility of L2W enables the inclusion of various important characteristics of PSEs. Future plans for L2W include the following features:

- *Database support for scenario and experiment management.* Using this feature, an L2W user will be able to define descriptions of new experiments by composing entries and representations from a relational database system (RDBMS). For example, the high-level specification introduced in the “Relation of PSE to DSS” section [Design simulations for the Roanoke watershed using the updated fisheries model from Marcus, but retaining the residential settlements of 1998] can be expressed as the SQL query

```
SELECT RunOff(*)
FROM Watersheds, Fisheries, Data
WHERE Fisheries.creator= “Marcus”
AND Watersheds.name= “Roanoke”
AND Data.type= “residential”
AND Data.year= 1998;
```

Query optimization algorithms built within the RDBMS can then help in determining an efficient way to answer the query based on past runs and the cost involved in conducting new simulations.

- *Data mining.* Massive quantities of data and graphs are generated from L2W simulations. Data mining can help harness this information by finding interesting, novel, and potential useful patterns. “Finding patterns” encompasses a variety of activities, from capturing regularities to deviation detection. In all of these cases, data mining of L2W simulation data can help one gain a better understanding of how to conduct simulations in the future. For example, assume that it is inferred that all experiments conducted using land cover obtained from source X cause substantial discrepancies in simulation results. Such a “pattern” could be an indicator for some special handling whenever data obtained from source X are to be incorporated. Results of data mining can also help in conducting sanity-checks and what-if analyses.
- *Computational steering.* In the L2W context, computational steering can help civil engineers, economists, and biologists interactively control the sequence of steps involved in integrating multiple models. Note that most model executions would require GIS support; if the economist is satisfied with the fidelity of the observations needed to run his or her model, then GIS services can be “steered” to support the modeling of other phenomena in greater fidelity. Computational steering thus recognizes the highly interactive and collaborative context in which simulations are conducted, and provides systems support to enhance problem solving.
- *Recommender systems.* Integrating more models into L2W increases the number of possible ways in which a simulation could be performed. Recommender systems are facilities that guide an L2W user in choosing an appropriate model (or model sequence) to achieve desired performance criteria. For example, such a system might suggest the following: “It is all right to use the low-resolution approximation for subwatershed 5, since it is known that model B2 is not very sensitive to changes involving impervious land.” A recommender system arrives at such recommendations, again, by data mining of a collection of past simulation data.
- *High performance computing.* The use of L2W for realistic case studies would require a lot of processing power, in terms

of both computing capabilities and storage-intensive operations. High performance computing addresses this larger context of doing computational science; models and codes are mapped onto supercomputers with multiple processors to achieve speedup in execution. Parallel computing is a well established area of computer science, and can help scale up simulations to problems and domains of realistic magnitude. With the use of a PSE such as L2W, there is an added advantage; such services are provided transparently and do not require that the domain scientist (hydrologist, economist, or biologist) have an understanding about the underlying system architecture or organization of data storage. A similar advantage applies in the university classroom where upper-level students knowledgeable in a discipline area are able to examine various land-use scenarios with only limited understanding of the many data layers required to execute a complex model such as HSPF. A working PSE can be an invaluable resource for engaging students on the subject of environmental and fiscal consequences of land development and urban sprawl.

## Conclusions

The attributes of a working PSE have been discussed at some length in this paper. It is gratifying that the National Research Council’s Committee on Watershed Management has developed a similar perspective on the need for integrated decision making tools to facilitate the assessment of impacts and their uncertainty. The committee, in its recent publication *New strategies for America’s watersheds* (1999), devotes two chapters to decision-making tools. It concludes with the thought that “effective watershed management requires integration of simulation models, data and expert judgement in a user-friendly decision support system to help decision-makers evaluate alternative approaches.” While a PSE is not mentioned per se by the committee, the document itself helps to substantiate a point of view advocated in this paper. It is also clear that the PSE concept presented in this paper embraces the primary attributes of a comprehensive planning model, and offers a glimpse into the future of watershed management. The technical challenge is to provide scientific linkage to the decision-making process.

## References

- Andreu, J., Capilla, J., and Sanchis, E. (1996). “AQUATOOL: A generalized decision-support system for water-resources planning and operational management.” *J. Hydrol.*, 177, 269–291.
- Artigas, F. J., Barrett, K. R., and Holowczak, R. (2001). “Digital meadowlands: a web-based decision support system for an urban, estuarine watershed.” *Proc., Integrated Decision-making for Watershed Management Symposium: Processes and Tools*, Virginia Water Research Center, Blacksburg, VWRRC P7-2001, 1–9.
- Berry, M. W., Flamm, R. O., Hazen, B. C., and MacIntyre, R. L. (1996). “The land-use change analysis system (LUCAS) for evaluating landscape management decisions.” *IEEE Comput. Sci. Eng.*, 3(1), 24–25.
- Bicknell, B. R., Imhoff, J. C., Kittle, J. L., Donigian, A. S. Jr., and Johnson, R. C. (1997). “Hydrological simulation program—FORTRAN (HSPF).” *User’s manual for version 11.0*, EPA National Exposure Research Laboratory, Research Triangle Park, N.C.
- Burrough, P. A. (1986). *Principles of geographical information systems for land resources assessment*, Oxford University Press, Oxford, U.K.
- CADSWES/Department of CEAE. (1993). “TVA environmental and river resources aid (TERRA) design document for 1993 prototype implementation.” Univ. of Colorado, Boulder.

- Carpenter, S., Brock, W., and Hanson, P. (1999). "Ecological and social dynamics in simple models of ecosystem management." *Conservation Ecology*, 3(2).
- Casanova, H., and Dongarra, J. (1997). "NetSolve: A network server for solving computational science problems." *Int. J. Supercomput. Appl. High Perf. Comput.*, 11(3), 212–223.
- Chen, C. W., Herr, J., Ziemelis, L., Goldstein, R. A., and Olmsted, L. (1999). "Decision support system for total maximum daily load." *J. Environ. Eng.*, 125(7), 653–659.
- Davis, J. R., Nanninga, P. M., Biggins, J., and Laut, P. (1991). "Prototype decision support for analyzing impact of catchment policies." *J. Water Resour. Plan. Manage.*, 117(4), 399–414.
- DeBarry, P. A., and Quimpo, R. G. (1999). *GIS modules and distributed models of the watershed*, ASCE, Reston, Va.
- Drashansky, T. T., Houstis, E. N., Ramakrishnan, N., and Rice, J. R. (1999). "Networked agents for scientific computing." *Commun. ACM*, 42(3), 48–54.
- Dunn, S. M., Mackay, R., Adams, R., and Ogleshorpe, D. R. (1996). "The hydrological component of the NELUP decision support system: An appraisal." *J. Hydrol.*, 177, 213–235.
- Fedra, K., and Jamieson, D. G. (1996). "The Waterware decision support system for river-basin planning: 3. Example applications." *J. Hydrol.*, 177, 199–211.
- Fortune, S. J., Gay, D. M., Kernighan, B. W., Landron, O., Valenzuela, R. A., and Wright, M. H. (1995). "WISE design of indoor wireless systems: Practical computation and optimization." *IEEE Comput. Sci. Eng.*, 2(1), 58–68.
- Foster, I., and Kesselman, C. (1997). "Globus: A metacomputing infrastructure toolkit." *Int. J. Supercomput. Appl.*, 11(2), 115–128.
- Gallopoulos, E., Houstis, E. N., and Rice, J. R. (1994). "Computer as thinker/door: Problem-solving environments for computational science." *IEEE Comput. Sci. Eng.*, 1(2), 11–23.
- Gannon, D., et al. (1998). "Component architectures for distributed scientific problem solving." *IEEE Comput. Sci. Eng.*, 5(2), 50–63.
- Goel, A., et al. (2001). "VizCraft: A problem solving environment for configuration design of a high speed civil transport." *Comput. Sci. Eng.*, 3, 56–66.
- Goel, A., Phanouriou, C., Kamke, F., Ribbens, C. J., Shaffer, C. A., and Watson, L. T. (1999). "WBCSim: A prototype problem solving environment for wood-based composites simulations." *Eng. Comput.*, 15(2), 198–210.
- Grimshaw, A. S., Ferrari, A., Lindahl, G., and Holcomb, K. (1998). "MetaSystems." *Commun. ACM*, 41(11), 46–55.
- Grimshaw, A. S., Weismann, J. B., West, E. A., and Loyot, Jr., E. C. (1994). "MetaSystems: An approach combining parallel processing and heterogeneous distributed computing systems." *J. Par. Distrib. Comput.*, 21(3), 257–270.
- Houstis, E. N., et al. (1998). "Parallel ELLPACK: A problem solving environment for PDE based applications on multicomputer platforms." *ACM Trans. Math. Softw.*, 24(1), 30–73.
- Ioannidis, Y., Livny, M., Gupta, S., and Ponnekanti, N. (1996). "ZOO: A desktop experiment management environment." *Proc., 22nd Int. VLDB Conf.*, T. M. Vijayaraman et al., eds., Morgan Kaufman, 274–285.
- Kant, E., Keller, R., and Steinberg, S., eds. (1992). "Working notes of the AAAI fall symposium on intelligent scientific computation." AAAI/MIT Press, Cambridge, Mass.
- Kittle, J. L. Jr., Lumb, A. M., Hummel, P. R., Duda, P. B., and Gray, M. H. (1998). "A tool for the generation and analysis of model simulation scenarios for watersheds (GenScn)." *Technical Rep. 98-4134*, U.S. Geological Survey, Menlo Park, Calif.
- Leavesley, G. H., Markstron, S. L., Brewer, M. S., and Viger, R. J. (1996). "The Modular Modeling System (MMS) — The physical process modeling component of a database centered decision support system for water and power management." *Water, Air, Soil, Pollut.*, 90, 303–311.
- Markus, S., Weerawarana, S., Houstis, E. N., and Rice, J. R. (1997). "Scientific computing via the Web: The net Pellpack PSE server." *IEEE Comput. Sci. Eng.*, 4(3), 43–51.
- New strategies for America's watersheds*. (1999). National Academy Press, Washington, D.C.
- Osmond, D. L., et al. (1997). "WATERSHEDSS: A decision support system for watershed scale nonpoint source water quality problems." *J. Amer. Water Resources Assoc.*, 33(2), 327–341.
- Pandey, S., Renee, G., Kyoung, J. L., Bernard, E., and Jon, H. (2000). "Impacts of land-use change: Information technology issues and a case study." *URISA J.*, 12(4), 5–21.
- Parker, S. G., Johnson, C. R., and Beazley, D. (1997). "Computational steering software systems and strategies." *IEEE Comput. Sci. Eng.*, 4(4), 50–59.
- Purtilo, J. M. (1994). "The POLYLITH software bus." *ACM Trans. Programming Languages and Systems*, 16(1), 151–174.
- Ram, S., Khatri, V., Hwang, Y., and Yool, S. (2000). "Semantic modeling and decision support in hydrology." *Photogramm. Eng. Remote Sens.*, 66(10), 1229–1240.
- Ramakrishnan, N., and Grama, A. Y. (1999). "Data mining: From serendipity to science." *IEEE Computer*, 32(8), 34–37.
- Ramakrishnan, N., and Grama, A. Y. (2001). "Mining scientific data." *Advances in computers*, M. Zelkowitz, ed., Vol. 55, Academic, San Diego, 119–169.
- Ramakrishnan, N., Houstis, E. N., and Rice, J. R. (1998). "Recommender systems for problem solving environments." *Technical Rep. WS-98-08*, H. Kautz, ed., AAAI/MIT Press, Cambridge, Mass., 91–95.
- Reitsma, R. F., Zagana, E. A., Chapra, S. C., and Strzepak, K. M. (1996). "Decision support systems (DSS) for water resources management." *Water resources handbook*, L. W. Mays, ed., McGraw-Hill, New York, 33.1–33.35.
- Rice, J. R., and Boisvert, R. F. (1985). *Solving elliptic problems using ELLPACK*, Springer, New York.
- Rice, J. R., and Boisvert, R. F. (1996). "From scientific software libraries to problem-solving environments." *IEEE Comput. Sci. Eng.*, 3(3), 44–53.
- Rubin, E. J., et al. (2000). "From landscapes to watersheds: A PSE for landuse change analysis." *Technical Rep. TR-00-09*, Dept. of Computer Science, Virginia Polytechnic Institute and State Univ., Blacksburg.
- Shaffer, C. A., Watson, L. T., Kafura, D. G., and Ramakrishnan, N. (2000). "Features of PSEs for computational science." *Proc., High Performance Computing Symposium, Advanced Simulation Technologies Conf. 2000*, A. Tentner, ed., Society for Modeling and Simulation Int., San Diego, 242–247.
- Simonovic, S. P., and Bender, M. J. (1996). "Collaborative planning support system: An approach for determining evaluation criteria." *J. Hydrol.*, 177, 237–251.
- Sol, H. G. (1983). "Processes and tools for decision support: Inferences for future developments." *Processes and tools for decision support*, H. G. Sol, ed., North-Holland, Amsterdam, The Netherlands, 1–6.
- Sugumaran, R., Davis, C., Meyer, J., Prato, T., and Fulcher, C. (2000). "Web-based decision support tool for floodplain management using high-resolution DEM." *Photogramm. Eng. Remote Sens.*, 66(10), 1261–1265.
- Voinov, A., and Costanza, R. (1999). "Watershed management and the web." *J. Environ. Manage.*, 56, 231–245.
- Weintraub, L. H. Z., et al. (2001). "Decision support system for Catawba River basin." *Proc., Integrated Decision-making for Watershed Management Symposium: Processes and Tools*, Virginia Water Resources Research Center, Blacksburg, VWRRC P7-2001, 890–905.
- Westervelt, J. D. (2001). *Simulation modeling for watershed management*, Springer, New York.