

---

**AWRA 2010 SPRING SPECIALTY CONFERENCE**  
**Orlando, FL**

**March 29-31, 2010**

**Copyright © 2010 AWRA**

---

HYDROSERVER: A PLATFORM FOR PUBLISHING SPACE-TIME HYDROLOGIC DATASETS

Jeffery S. Horsburgh, David G. Tarboton, Kimberly A. T. Schreuders,  
David R. Maidment, Ilya Zaslavsky, and David Valentine\*

**ABSTRACT:** The Consortium of Universities for the Advancement of Hydrologic Science, Inc. (CUAHSI) is developing a Hydrologic Information System (HIS) that supports sharing of distributed hydrologic data through Web services and tools for data discovery, access, and publication. HydroServer is a computer server that contains a collection of databases, Web services, and software that allows data producers to store, publish, and analyze space-time hydrologic datasets. HydroServer is designed to permit local control of the data, while still being part of a distributed, national/international system allowing universal access to the data. We describe the HydroServer architecture and software stack, including tools for managing and publishing time series data for fixed point monitoring sites, as well as spatially distributed, GIS datasets that describe a particular study area, watershed, or region. We then describe a case study implementation of HydroServer for publishing water resources related space-time datasets for the Inland Northwest Research Alliance (INRA) Constellation of Experimental Watersheds (ICEWATER) network of research watersheds in the northwestern United States.

**KEY TERMS:** Hydrologic Information Systems, CUAHSI, Cyberinfrastructure

INTRODUCTION

HydroServer is part of the CUAHSI HIS (Maidment, 2009; Horsburgh et al., 2009; Tarboton et al., 2009) and was designed to enable producers of hydrologic data to join a growing, distributed network of published water data services by establishing their own server and publishing their own data. HydroServer software makes data available via Web services on the Internet, and, when services are registered with a central cataloging service called HIS Central, the data become discoverable through HIS client applications. HydroServer includes software for management of data on the server, as well as Internet applications for presenting data to consumers. HydroServer also enables publication of spatial datasets for a study area or region and the linkage of spatial datasets with point observational datasets that are hosted on the same server.

This paper describes the second generation of HydroServer software. Initially developed in 2007 under the name “HIS Server” and deployed at several Hydrologic Observatory Test Beds in support of the water data publication workflow (Whitenack et al., 2007a; 2007b), HydroServer existed in two versions: HIS Server and HIS Server Lite (the latter was built from freely available components, while the former expected the hosting organization to provide Microsoft SQL Server and ArcGIS Server licenses). Of the 44 water data services registered at HIS Central, 38 are now served by such HydroServers installed at universities and partner projects. The new HydroServer described in this paper is an enhanced version. Standardizing on MS SQL Server 2008 and ArcGIS Server 9.3.1, it adds a new mapping interface, Capabilities database and services, Time Series Analyst, and a configurable HydroServer website.

HydroServer uses the concept of a study area, or “region,” as an area within which hydrologic data are collected. A region may be an experimental watershed, field site, or other geographic space. The region concept is important as an organizing principle within HydroServer as it provides a way for spatial datasets describing a region to be associated with point observations data collected within that region. It also enables HydroServer administrators to host data for several regions on a single HydroServer. Each region can be associated with one or more observational data services (e.g., point observations from stream gages, weather stations, etc.) as well as one or more spatial data services that describe the region (e.g., a watershed boundary, stream locations, digital elevation model, etc.). Thus, HydroServer is a platform for a Digital Watershed that describes both the spatial and temporal characteristics of a region.

---

\* Respectively, Research Assistant Professor, Utah Water Research Laboratory, Utah State University, 8200 Old Main Hill, Logan, UT 84322-8200, (435) 797-2946, [jeff.horsburgh@usu.edu](mailto:jeff.horsburgh@usu.edu); Professor, Utah Water Research Laboratory, Utah State University; System Administrator, Utah Water Research Laboratory, Utah State University; Professor, Center for Research in Water Resources, University of Texas at Austin; Director, Spatial Information Systems Laboratory, San Diego Supercomputer Center; Research Programmer, Spatial Information Systems Laboratory, San Diego Supercomputer Center

## HYDROSERVER FUNCTIONAL COMPONENTS

Figure 1 shows an overview of the major components of HydroServer. Observations made at fixed points are loaded into one or more Observations Data Model (ODM) databases (Horsburgh et al., 2008). Each ODM database is attached to a WaterOneFlow Web service that publishes the data on the Internet in Water Markup Language (WaterML) (Zaslavsky et al., 2007). Spatial datasets are published as services using ESRI's ArcGIS Server. A Capabilities database contains a listing of all services that have been published (both WaterOneFlow and spatial services), as well as appropriate metadata to describe each service. A configuration tool enables HydroServer administrators to edit the contents of the Capabilities database.

A suite of Web applications connects to the Capabilities database and makes the data published on the HydroServer available via a Web browser. These include: 1) an Internet map that combines GIS datasets and time series published using ODM and WaterOneFlow services for a region; 2) the Time Series Analyst for visualizing observational data, 3) a HydroServer website that describes the regions and services published on the HydroServer; and 4) the Capabilities Web service, which publishes the capabilities (i.e., a metadata description of the services that the HydroServer contains) of the HydroServer in extensible markup language (XML) format. The following sections describe the components of HydroServer, the latest versions of which are available at <http://his.cuahsi.org>.

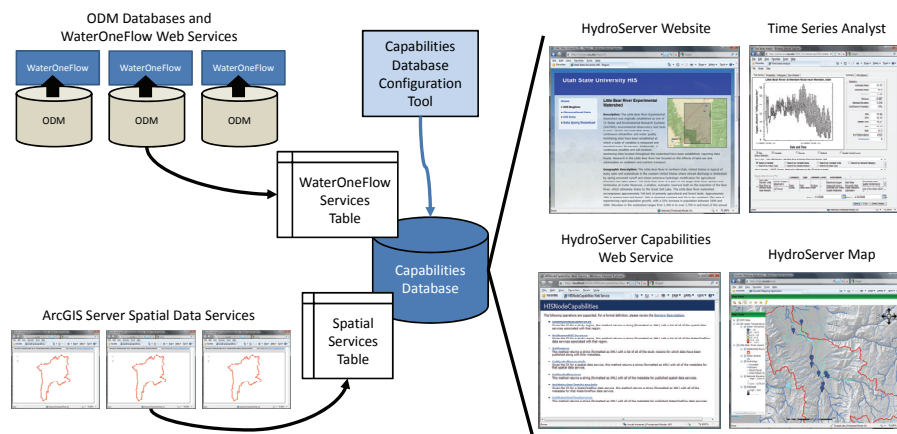


Figure 1. General Overview of HydroServer Components.

## Publishing Point Observations with ODM and WaterOneFlow

ODM and the WaterOneFlow Web services serve as the basis for publishing point observations on a HydroServer. Point observations are loaded into one or more ODM databases, each of which is then connected to a WaterOneFlow Web service. The WaterOneFlow Web services transmit the data over the Internet in WaterML format. Additionally, HydroServer includes several utilities for working with ODM databases. ODM, WaterOneFlow, and the ODM utilities are described briefly below.

ODM is a relational model for storing, managing, and manipulating point observations data. Design specifications for ODM and a schema diagram can be found at <http://his.cuahsi.org>. Each HydroServer implements one or more ODM databases, depending on the number of research watersheds/regions hosted by the HydroServer. All of the point observations data for a region are entered into one or more ODM databases. ODM provides a persistent storage mechanism for the data. Because ODM databases are implemented in a relational database management system (e.g., Microsoft SQL Server), they support a variety of applications through the use of SQL queries that can be passed to the database to retrieve and manipulate data. Additionally, once data are in an ODM database, they can be published using the WaterOneFlow Web services.

The ODM Data Loader (ODMDL) loads tabular data (Microsoft Excel or comma separated values .csv) into ODM. ODMDL validates the data as it is loaded, thereby protecting the security and consistency of an ODM database and minimizing the potential for human caused errors. The ODM Streaming Data Loader (ODMSDL) was designed for loading continuously measured sensor data into an ODM database. Similar to ODMDL, ODMSDL is a file based data loader, but takes as input datalogger files that have a single date/time column and one or more columns of data. ODMSDL provides tools for mapping streaming data files to the ODM schema and for specifying required metadata. ODMSDL can be run manually or automatically as a scheduled task. ODMDL and ODMSDL enable the loading of data that are fully qualified with valid metadata into ODM without the need for specialized programming.

ODM Tools (Figure 2) was created for managing data stored in an ODM database. ODM Tools automates many of the most common database transactions and allows users to interact with data in an ODM database through a graphical user interface. ODM Tools is organized into three general areas: 1) query and export; 2) visualize; and 3) edit. The "Query and

Export” functionality allows users to find the data that they are interested in and export it to a simple format that can be used with a variety of analysis software. This allows users to work with copies of data exported from their ODM database while preserving the provenance of the data via an accompanying metadata file. The “Visualize” functionality allows quick plotting and summarization of data using a variety of plot types and statistics. The “Edit” capability provides functionality for editing data series (e.g., for performing QA/QC on continuous data streams) and for creating new data series from existing data series. The source code of ODM Tools provides an example of how software can be built on top of an ODM database.

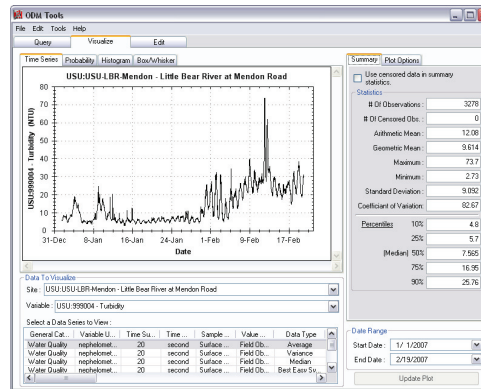


Figure 2. ODM Tools Graphical User Interface.

The WaterOneFlow Web services are used to publish data stored in an ODM database on the Internet. WaterOneFlow Web services transmit data in WaterML format and consist of a set of methods (e.g., GetSites, GetSiteInfo, GetVariableInfo, GetValues) that can be called from many different programming languages and software environments for retrieving data from an ODM database over the Internet. These methods have been implemented within a single Web application that is easily installed and configured. Each HydroServer can implement WaterOneFlow Web services for each ODM database that contains data to be published. This ensures that all of the data published on HydroServers are available on the Internet in a standard, interoperable (i.e., platform and programming language agnostic) format. In addition, WaterOneFlow services can be registered with CUAHSI HIS Central (<http://hiscentral.cuahsi.org/>), enabling the services to be discovered and more easily accessed by HIS client applications like HydroDesktop (<http://www.hydrodesktop.org>) (Ames et al., 2009).

### Publishing Spatial Datasets Using ArcGIS Server

HydroServer uses ArcGIS Server to publish GIS datasets as Web services using Open Geospatial Consortium (OGC) standards (OGC, 2004; 2005; 2008). Vector data are published using the Web Feature Service (WFS) standard and raster data using the Web Coverage Service (WCS) standard. Hybrid services containing both vector and raster data can be published as image services using the Web Map Service (WMS) standard. OGC standards are open formats that can be consumed by several GIS software clients. GIS datasets are inherently tied to a location or spatial extent, and it is for this reason that the concept of a study area or “region” is so important within HydroServer. The region gives context to the geographic extent that the spatial dataset describes. For example, a user might create a watershed boundary for the Little Bear River Experimental Watershed region or a digital elevation model for the T.W. Daniels Experimental Forest region. A region may have many spatial datasets that describe its characteristics.

Spatial data services can contain a single spatial dataset such as a watershed boundary or digital elevation model, or multiple spatial datasets can be combined to create a composite map service. This provides flexibility in defining spatial data services that are appropriate for a region. Spatial data services can then be reused for multiple regions – for example, a single map service that covers the whole United States might be used as the base map for multiple regions for which data are published on the HydroServer.

### Publishing HydroServer Capabilities

Each HydroServer has a Capabilities database (Figure 3) that catalogs the regions for which data have been published and the list of services that have been published on the server. The database also contains metadata describing each region and service. A configuration tool has been built for defining the contents of this database. It provides an interface for defining regions, services, and the metadata describing each. Inputs to the forms are written to the appropriate tables in the database.

Each HydroServer can then implement Web services that publish its capabilities. The Capabilities Web service includes methods that return, in XML format, the list of regions for which data have been published, the published point observations

data services, and the list of published spatial data services, along with appropriate metadata for each. By doing so, all of the capabilities of the HydroServer are published in an XML format that can be discovered by central registration and cataloging services (e.g., HIS Central). The Capabilities Web service makes a HydroServer self describing. A programmer can discover all of the capabilities of the HydroServer simply by calling the Capabilities Web service, and metadata can be registered and harvested automatically. The Capabilities Web service connects directly to the Capabilities database and uses the tables that list the services that have been published on the HydroServer. As HydroServer administrators define new regions or services within their HydroServer database, they are immediately available to the Capabilities Web service. The service has the following methods: 1) GetRegions; 2) GetRegionInfo; 3) GetWaterOneFlowServices; 4) GetWaterOneFlowServiceInfo; 5) GetMapServices; 6) GetMapServiceInfo; 7) GetRegionWaterOneFlowServices; and 8) GetRegionMapServices.

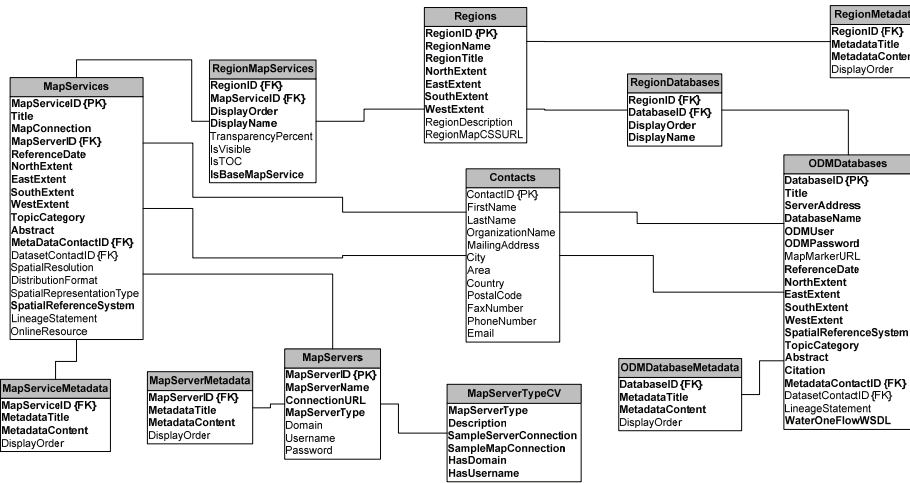


Figure 3. Database Schema for the HydroServer Capabilities Database.

### HydroServer Data Presentation and Visualization Tools

A suite of data presentation and visualization tools has been created for HydroServer. The suite includes the HydroServer Website, the Time Series Analyst, and the HydroServer Internet map application and is intended to provide a public graphical user interface to the data holdings of the HydroServer.

Each HydroServer can implement an Internet map interface to the observational and GIS data services that are published on that server (Figure 4). The map is a Web application that runs in a browser and allows users to browse data that are published for a particular region. Each region can have its own map. The application is based on ArcGIS Server, and combines both spatial datasets and the locations of point observations in a single map (e.g., it plots the locations of monitoring sites and provides site information and links to the data when users click on a site on the map). The map application dynamically generates its content using information contained within the Capabilities database. It is capable of presenting data from multiple published observational and GIS data services for a region, and the list of services displayed for a region can be edited by HIS Sever administrators. When services are added to the Capabilities database, they automatically appear in the map application. An example implementation of the map can be seen at <http://maps.usu.edu/map/>.

Each HydroServer can host an instance of the Time Series Analyst (Figure 4), which enables data consumers to visualize and generate descriptive statistics for selected point observations datasets. The Time Series Analyst is a Web application that enables screening/preview of datasets prior to download so that users can make sure that the data are what they want. An example of the Time Series Analyst is available at [http://tsa.usu.edu/odm\\_tsa/](http://tsa.usu.edu/odm_tsa/). The Time Series Analyst can be linked to the Internet map application so that when users click on a monitoring site on the map they are presented with a link to visualize/summarize the available data at the selected site using the Time Series Analyst. The Time Series Analyst is capable of connecting to any number of ODM databases hosted on a HydroServer, and has a unique URL calling interface that can be used to launch the application with a selected ODM database, site, variable, and date range.

Each HydroServer can also implement a website that provides information about the HydroServer and details for each region and published service that reside on that HydroServer (see <http://icewater.usu.edu>). The majority of the content of this website is dynamically generated from the Capabilities database and the ODM database(s) that hold the observational data. This enables the website to be dynamically updated through adding content to the Capabilities database rather than editing the code of the website. The HydroServer website contains the following components/pages: 1) an overall/opening page that describes the HydroServer and descriptive information about each region for which data is published on the server; 2) a dynamically generated listing of published point observations data services; 3) a dynamically generated listing of published

spatial data services; 4) links to the map application and Time Series Analyst; and 5) a data query and download page that allows users to more easily query for data from one or more published observational data service and then download the data.

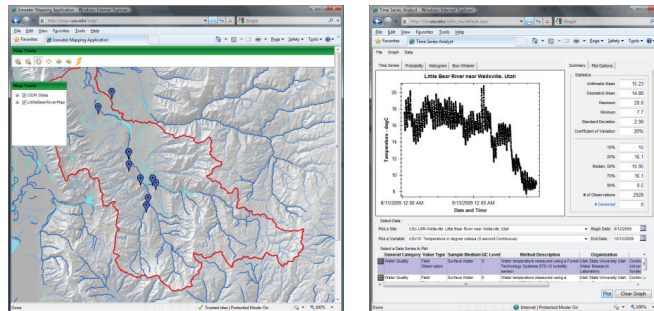


Figure 4. HydroServer Map Application and Time Series Analyst for the Little Bear River Experimental Watershed Region.

### ICEWATER: USING HYDROSERVER TO BUILD A REGIONAL NETWORK OF PUBLISHED DATA

Over the past several years, researchers at universities affiliated with the Inland Northwest Research Alliance (INRA) have been collecting water resources datasets at a number of experimental watersheds in the western United States. Experimental watersheds in the INRA region span a number of climate, human development, and disturbance gradients, and researchers are investigating several different research themes, including snowmelt responses to climate change, groundwater-surface water interactions, modeling of hydrologic response, land use change, and arctic river processes. It is hoped that integration of data from these watersheds will facilitate cross-site comparisons and larger scale studies that synthesize information from diverse settings, making the network as a whole greater than the sum of its parts.

The INRA Water Resources Consortium Constellation of Experimental Watersheds Information System Network (ICEWATER) is a distributed network of HydroServers that is being used to publish and integrate the data holdings from each INRA university. The goals of ICEWATER are: 1) establishment of a common information system for data sharing, analysis and archiving that builds upon and extends the CUAHSI HIS; 2) establishment of a common modeling framework to facilitate sharing and model interoperability; and 3) establishment of common base characterization datasets such as digital elevation models (DEMs) from LIDAR, land cover and land use from remote sensing, for research sites within the ICEWATER network that provide detail beyond nationally available information.

ICEWATER is an ongoing effort that is using the CUAHSI HIS software to build a regional HIS that is also part of the national network of data published using HIS tools. A HydroServer has been established at each INRA university, and on each HydroServer the data holdings of one or more regions (e.g., experimental watersheds) are being published. Each ICEWATER HydroServer is in the process of implementing one or more ODM databases connected to WaterOneFlow Web services, spatial data services published using ArcGIS Server, a Capabilities database and Capabilities Web service, and the whole suite of data access, presentation, and visualization tools described above. As WaterOneFlow services are coming online, they are being registered with HIS Central so that they become part of the national HIS data network.

One additional component of ICEWATER is functionality that is referred to as ICEWATER Central (<http://icewater.inra.org>), where support is provided for HydroServer administrators and data managers at each of the INRA Universities. ICEWATER Central consists of personnel who provide support via telephone and via a group email listserv, as well as a website with the following components: 1) general information about INRA and ICEWATER; 2) resources for HydroServer administrators and data managers, including a listing of all hardware and software needed to establish a HydroServer; and 3) a catalog of and presentation of the published data services from each ICEWATER HydroServer.

The ICEWATER Central functionality is not intended to replace functionality of the CUAHSI HIS Central website. Rather, it is intended to give the ICEWATER Network an identity that is separate from the CUAHSI HIS, but in a way that is still entirely compatible with the national network of published data that is enabled by the CUAHSI HIS. Because each ICEWATER data manager is registering their services with HIS Central, the entire contents of the ICEWATER regional HIS will be available through HIS Central and can be discovered and accessed using HIS tools like HydroDesktop.

### SUMMARY

In this paper, we have described the CUAHSI HydroServer, which is a computer server that integrates collections of point observations and spatially distributed GIS datasets, Web services, and software that enables data producers to store,

publish, and analyze space-time hydrologic datasets. HydroServer was designed to permit data managers affiliated with universities, research watersheds, experimental sites, etc. to establish their own server and publish their own data resources. This permits local control of the data, and through registration with HIS Central, each HydroServer can become part of a distributed, national/international system that enables universal access to the data.

The HydroServer software stack includes ODM, the WaterOneFlow Web services, ESRI's ArcGIS Server for publishing spatial datasets, a Capabilities database and Web service, and a suite of data presentation and visualization tools that includes a map application, Time Series Analyst, and the HydroServer website. Also included are a set of utilities for managing and working with time series data for fixed point monitoring sites in ODM, including ODM Tools, ODMDL, and ODMSDL.

We have also described how HydroServer is being used to build a regional HIS that has both a regional identity and participates in the national network of published hydrologic data that is enabled by the CUAHSI HIS. Indeed, HydroServer is the basis for the INRA ICEWATER network where a HydroServer has been established at each INRA university. This has allowed data for a variety of experimental watersheds within the western United States to now come online, whereas before they were confined to the files of individual researchers or available on the Internet in inconsistent formats. It is anticipated that the growing, distributed network of HydroServers, such as the ones being built within the ICEWATER network, will facilitate cross-site comparisons and larger scale studies that synthesize information from diverse settings, making the network as a whole greater than the sum of its parts in advancing hydrologic research.

#### ACKNOWLEDGMENTS

This work was supported by National Science Foundation grants EAR 0413265 and EAR 0622374 and by the Inland Northwest Research Alliance (<http://www.inra.org>). Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation or INRA.

#### REFERENCES

- Ames, D. P., J. S. Horsburgh, J. Goodall, D. G. Tarboton, T. Whiteaker, and D. R. Maidment (2009), Introducing the open source CUAHSI Hydrologic Information System desktop application (HIS Desktop), in Anderssen, R. S., R. D. Braddock, and L.T.H. Newham (eds) 18th World IMACS Congress and MODSIM09 International Congress on Modelling and Simulation, Modelling and Simulation Society of Australia and New Zealand and International Association for Mathematics and Computers in Simulation, July 2009, pp. 4353-4359, ISBN: 978-0-9758400-7-8.
- Horsburgh, J. S., D. G. Tarboton, D. R. Maidment, and I. Zaslavsky, 2008. A relational model for environmental and water resources data, *Water Resources Research*, 44, W05406, doi:10.1029/2007WR006392.
- Horsburgh, J. S., D. G. Tarboton, M. Piasecki, D. R. Maidment, I. Zaslavsky, D. Valentine, and T. Whitenack, 2009. An integrated system for publishing environmental observations data, *Environmental Modeling and Software*, 24, 879-888, doi:10.1016/j.envsoft.2009.01.002.
- Maidment, D. R., ed., 2009. CUAHSI Hydrologic Information System: 2009 Status Report, Consortium of Universities for the Advancement of Hydrologic Science, Inc, 79 p., [http://his.cuahsi.org/documents/HISOverview\\_2009.pdf](http://his.cuahsi.org/documents/HISOverview_2009.pdf).
- Open Geospatial Consortium, 2004. OGC Web Map Service Interface, Version 1.3.0, OGC 03-109r1, [http://portal.opengeospatial.org/files/?artifact\\_id=4756](http://portal.opengeospatial.org/files/?artifact_id=4756).
- Open Geospatial Consortium, 2005. Web Feature Service Implementation Specification, Version 1.1.0, OGC 04-094, [http://portal.opengeospatial.org/files/?artifact\\_id=8339](http://portal.opengeospatial.org/files/?artifact_id=8339).
- Open Geospatial Consortium, 2008. Web Coverage Service (WCS) Implementation Standard, Version 1.1.2, OGC 07-067r5, [http://portal.opengeospatial.org/files/?artifact\\_id=27297](http://portal.opengeospatial.org/files/?artifact_id=27297).
- Tarboton, D. G., J. S. Horsburgh, D. R. Maidment, T. Whiteaker, I. Zaslavsky, M. Piasecki, J. Goodall, D. Valentine, and T. Whitenack, 2009. Development of a community Hydrologic Information System, in Anderssen, R. S., R. D. Braddock, and L.T.H. Newham (eds) 18th World IMACS Congress and MODSIM09 International Congress on Modelling and Simulation, Modelling and Simulation Society of Australia and New Zealand and International Association for Mathematics and Computers in Simulation, July 2009, pp. 988-994, ISBN: 978-0-9758400-7-8.
- Whitenack, T., D. Valentine, I. Zaslavsky, D. Djokic, 2007a. Hydrologic information system server: the software stack and the initial deployment experience. *Geoinformatics 2007*, [http://gsa.confex.com/gsa/2007GE/finalprogram/abstract\\_122332.htm](http://gsa.confex.com/gsa/2007GE/finalprogram/abstract_122332.htm).
- Whitenack, T., Zaslavsky, I., Valentine, D., Djokic, D., 2007b. Data Access System for Hydrology, *Eos Trans. AGU*, 88(52), Fall Meet. Suppl., Abstract H13H-1685.
- Zaslavsky, I., Valentine, D., Whiteaker, T. (Eds.), 2007. CUAHSI WaterML. OGC Discussion Paper OGC 07-041r1. Version 0.3.0. <http://www.opengeospatial.org/standards/dp>.