

AN ADVANCED SYSTEM FOR THE PRODUCTION AND VIEWING OF WEATHER PRODUCTS ON THE WORLD-WIDE WEB

David P. Wojtowicz*, Steven E. Hall, Joel P. Plutchak, Scott L. Olthoff,
Edward A. Mlodzik, Paul S. Dekker, Daniel J. Bramer, Mohan K. Ramamurthy

University of Illinois at Urbana-Champaign
Department of Atmospheric Sciences

1. BACKGROUND

Most of the best known and most extensive of the WWW servers that offer weather information products started as a single page listing and expanded from there to the detailed multi-page collections that they are today. Such was the case with the sections of the University of Illinois at Urbana-Champaign Department of Atmospheric Sciences' servers. Excluding the sections that offered educational material and the widely successful interactive weather map tools, the rest of the sections offering the bulk of the weather products however, unfortunately fell into disarray...victims of uncoordinated growth and the ultimately unmanageable underlying structures that supported them.

So armed with experiences, successes, failures, and lots of wishlist items, we started out with a clean slate and undertook the design of a completely new weather server in order to "do it right this time."

The following details the design process and goals of the *Weather World 2010*TM server project. (Because much of the implementation of this design was still in progress (and flux) at the time of this writing, details of the actual implementation will be presented at a later date.)

2. INITIAL DESIGN PROCESS

The first phase of the design process was simply to spend some time dreaming up the kind of server we would like to have. To accomplish this, we enlisted the help of several weather and web enthusiasts within our department, and over the course of a month or so, spent time exploring dozens of existing sites, comparing and contrasting features, pointing out the good and the bad. We also listed all the things (including both products and technologies) that could be rather useful to have, but were not yet offered by other WWW-based weather servers.

In looking at all of the creative things people were doing on the web in many other fields and at all the new WWW technology that has recently become available, we ultimately knew we both desired and required something that went well beyond a big list of nice looking products. For example, the new server should take advantage of technologies such as JavaTM¹ to offer interactive product viewers to the user. It should take advantage of new server technologies that make it possible to customize the site on a per-user basis for example. It should also be simple to maintain. These concepts will be explored in more detail in the following sections.

3. USER INTERFACE DESIGN

The following detail some of the user-visible design elements of *Weather World 2010*:

3.1 Organization and Navigation

A well thought-out, hierarchical organization and navigation system to move around are keys to unlocking the full potential of a web site for a user. One has to achieve a balance between a hierarchy that is too "flat" (in which one or two levels contain too many items forcing the user to scroll up and down lengthy pages looking for the items they want) and one that is too deep (with many levels that the user must descend and climb back out of to move from one area to another).

Considerable attention was given to this issue early in the design process. Dozens of HTML pages were mocked up to explore different layout styles and menu systems. Through an iterative process of constructive criticism and reworking we sorted the good designs from those that were "too crowded", "hard to understand", "required too much scrolling" to arrive at the present design.

The WW2010TM navigation system consists of a series of vertically arranged menus on the left edge of the page. Every page has an "Available Resources" menu that lists the most general categories of products

* *Corresponding Author Address:* David Wojtowicz, University of Illinois at Urbana-Champaign, Dept. of Atmospheric Sciences, 105 S. Gregory St., Urbana, IL 61801 e-mail: davidw@uiuc.edu

¹ JavaTM is a trademark of Sun Microsystems Inc. (see <http://java.sun.com/>)

available. The menu item specifying the current category is highlighted. Submenus below the resources menu offers access to subcategories. By clicking on a category and a subcategory, one can go to any major section of the server from any other section with only one or two clicks. Additional menus appear on each page that control user interface mode (see below), regional views, and other settings. These retain their settings from page to page.

3.2 Levels of Presentation

While we pursued WWW page designs featuring many inline graphics and icons as is now commonplace, we continued to get mail from users of our old gopher-based server lamenting its demise. They liked the speed of the simple text-based menus. This is especially important to users with slow network connections and even well connected users who wish to move quickly through pages without waiting for icons to load.

In recognition of this, we offer several “user interface” modes, selectable via a small menu that appears in the menu section of each page. The default “high graphics” mode offers all of the fancy icons that many users have come to expect. The “low graphics” mode strips away all extraneous graphics leaving only fast loading text in nearly the same format. In low graphics mode, the only graphics provided are the maps and images that the user specifically requests.

To further accommodate various levels of users, a number of the products are offered in multiple sizes to serve both the user with modem and a 13” monitor and the user with a direct Ethernet connection and a 21” monitor.

In addition to the high and low graphics modes there are two other presentation modes. One makes use of HTML forms-based pulldown menus to very compactly offer access to dozens of products on a single page. The other enables Java-based interactive capabilities as described below.

3.3 Interactive Viewers

When interactive viewing mode is enabled by the user, products are no longer delivered as simple images to only be displayed by the browser or helper application then dismissed. Instead, a Java-based interactive viewer would automatically load and pop up to display the product. For most kinds of products this viewer would offer features such as the capability of animating through a time-series of images. It might also display the exact latitude and longitude of the point under the cursor. In addition to the standard features, additional product specific features would be enabled automatically.

Point data extraction is one such feature. For example, in the satellite image viewer, when one moves the cursor over an infrared image, a numerical readout in the margin would display the derived cloud top temperature at that point. On a surface temperature contour map, the specific temperature at that point (subject to interpolation) might be displayed. On satellite images, it would also be possible to select from a number of different color enhancement tables.

These are just a few of the interactive features that would be made available in this way. They require no additional software or plug-ins on the user’s end, only a Java-capable WWW browser. A prototype tool that demonstrates similar features is the closely related NASA-funded HORIZON Project’s “Weather Visualizer”² portions of which will be adapted for this purpose. (Plutchak, 1997)

3.4 Everything Explained

The most frequent type of e-mail that we get in response to our existing servers comes from those who wonder about things like what a certain weather map symbol means, what the colors in an enhanced satellite image represent, or more generally, how to thunderstorms form and what is meant by the terms like “vorticity advection.”

Since the beginning, it has been our goal to make WW2010 as much an educational server as an information server. To this end, every single weather product will have a descriptive page hyperlinked to it. This page will detail the specific features of the particular product such as a radar map, satellite image, etc. and describe what it is useful for. Symbols and color codes will be explained. The fields displayed by the map such as surface pressure contours will also be briefly described. Examples of significant features indicated by a map or image of this type will also be provided.

More extensive discussions on topics such as the concept of vorticity, the structure of hurricanes, and the characteristics of cold and warm fronts will be provided throughout by integrated links into the ever growing collection of WWW-based multimedia instructional modules on meteorological topics developed as part of the closely-related, NSF-funded CoVis project³ (Hall, 1997)

3.5 Archived Data

Significant weather events often bear further study. On many servers, however, the data from yesterday’s big storm is quickly purged within 24 hours. There’s no

² <http://covis.atmos.uiuc.edu/java>

³ <http://covis.atmos.uiuc.edu/guide>

going back to study it in detail, say in a classroom setting, the next day.

In WW2010 a month or more of easily accessible archive files for many of our products will be retained online. In addition, selected data from significant weather events such as notable hurricanes, blizzards, and other severe weather outbreaks will be kept online in readily available special case-study collections. These collections may offer additional material such descriptive text, special charts, etc.

3.6 Text Products

Though often omitted from many WWW-based weather servers, perhaps because they are unglamorous, there are a great many useful text-based products available from the standard DDPLUS data service. At one time we offered many of these products on our gopher-based Weather Machine server. Despite the fact that this server has been out of operation in this capacity since March of 1996, we continue to get several pieces of e-mail each week asking when these products are going to return as many people had found them to be quite useful. For this reason, we will endeavor to make an extensive collection of these products available in WW2010.

3.6 Personalized Pages

A number of commercial web sites that offer news and information now offer the capability to display a personalize page for the user. This page displays a customized set of favorite or frequently checked products and links selected by the user from the server's other sections. In the case of WW2010, similar capabilities will exist to allow a user to display a page, containing in one place, only a list of products that they find most useful.

4. UNDERLYING SERVER DESIGN

The end user only sees a small part of the whole server. A large portion of the server's work goes on behind the scenes. This ranges from the process of turning raw data into the maps and images to the process of delivering the HTML and GIF files to the user and everything that happens in-between.

4.1 Dynamic HTML

Hundreds of lines of complex HTML markup go into each page of WW2010. Multiply this by dozens of pages in multiple versions (low graphics vs. high graphics) and one ends up with a large collection of HTML files that are difficult to maintain by hand. Additionally, many of the pages need to be continuously updated to reflect current product listings as new products arrive each hour.

The answer, of course, is to dynamically generate the HTML content of each page as it is requested rather than simply serve static documents. This is the approach that has been adopted by a number of commercial sites, especially those that offer online shopping. Product lists from the business's inventory system are automatically converted to web page catalogs, the format of which is based on a template definition. In this way the web pages always reflect the current product inventory. Because the complex parts of the page formatting are defined centrally by a template, the unique content of each page is all that's contained within each page's file. This simplifies the construction and editing of the page. It also makes it possible to instantly change the layout of all the pages by simply modifying the template.

This approach has significant other benefits. By having multiple templates available, multiple versions of the same page can be generated. In the case of WW2010, an example would be the high and low graphics versions of the pages. Conditional statements within the templates can also change the formatting based on special conditions, such as the detection of a browser that does or does not support certain special features.

The most interesting aspect of this approach is the ability to personalize the contents of the server for each user. This has been made possible by the advent of what are called "cookies". A cookie is essentially a unique id number assigned to a particular browser instance running on a particular client machine. This allows a server to differentiate between many simultaneous browsing sessions. Once the server can differentiate between sessions, it can store specific information for each session that persists across multiple page accesses and even across multiple return visits to the server. This works without having to necessarily specify a user name or password and is supported now by most major browsers.

This makes it possible to do things like remember what subcategory of each category the user was last in and return to it when that category is selected. In WW2010, it will also for example, allow the server to automatically default to low or high graphics modes when the user returns to the site at a later time depending on the setting they used the last time. It would also allow the server to offer a personal page of favorite products as described above.

4.2 Product Generation Management

One of the more complex aspects of operating a WWW-based weather product server is making sure all the products are generated and made available in a timely manner. Each needs to be updated each time the raw data it was derived from is updated. Typically this

is accomplished by a series of scripts that call the map and image making programs, etc. These scripts are generally launched a specific times by a mechanism such as the UNIX cron system.

While this sort of system works well, there are limitations. By processing things at specific times it may delay the availability of a product when the raw data has arrived earlier or may fail to process it at all if the data has arrived later than usual. In WW2010, we will ultimately employ an event driven product management system to deal with this issue. The arrival of new data will trigger the generation of products that depend on it and in turn any products that depend on an intermediate product (such as a converted data file). This will deal equally well with both current data and raw archival data inserted into the system at a later time. It will also later generate any products that were missed during any downtime of the processing system.

A relational database will be used to track both product/dataset dependencies and the collections of datasets and products currently in the system. This information will be accessed by the dynamic page generating mechanisms described above to accurately display listings of the current and archived products that are available within the system at any given time along with their locations.⁴

Because the generation of many products can require extensive computational resources, it is often both necessary and desirable to divide this workload among multiple machines. Typically this is done by hand in an ad-hoc fashion. If there are many products, it tends to become difficult to remember which machine does what and from what directory it does it. This will be handled by a distributed queuing system in WW2010.

When the WW2010 product management system determines that a product needs to be generated, it will submit a job to DQS (SCRI's Distributed Queuing System)^{5,6} DQS in turn, finds an idle machine in an easily reconfigurable collection of several with the basic capabilities needed for the specific task and executes the job there. NFS-mounted directories provide filesystem connectivity between the systems. DQS had been in use for pre-WW2010 product generation at UIUC for over a year.

⁴ A prototype system called UPRO implemented (in limited fashion) both the event driven product generation and database tracking functions described here. It was responsible for the production most of the products in the now defunct original UIUC Weather World web-site.

⁵ SCRI: <http://www.scri.fsu.edu/>

⁶ DQS: <http://www.scri.fsu.edu/~pasko/dqs.html>

Finally, one additional design goal is to further centralize the product generation by creating a small number of complex product generation scripts that read hierarchical configuration settings and can produce multiple products rather than have many small scripts, one for each product.

5. A GROUP EFFORT

No matter how much technology one employs, human input remains an absolutely essential ingredient. Throughout the design and implementation of WW2010, we sought to involve both as many of those within our department who were interested in the project as well as several college and high school students at other institutions in cooperation with the CoVis project. We did this with the full support of our department head and several faculty members.

In addition to the initial brainstorming and periodic review sessions we had with several interested individuals, during the early stages we also posted charts in the hallway of our building listing the scope and kinds of products that would be available. We invited passers-by to mark them up with comments and additions, providing additional input.

For example, one student commented, "Why stop at the tropopause?" indicating her research interest in events higher up and possibly resulting in the future addition of periodically updated stratospheric ozone concentration maps... something we had not thought of, but likely of much interest to the public.

Once the design was farther along we formally presented our plans to our graduate students at a seminar and invited them all to become contributors to the project. We were particularly seeking input in the form of text and graphics for the many informational pages that either briefly described individual products or explained topics in detail, especially in their specific areas of the student's interest and expertise.

To facilitate this we offered workshops on HTML and web design along with individual instruction as needed. This benefited the students by providing them with both skills and experience in the area of web authoring.

6. ADDITIONAL INFORMATION

The *Weather World 2010* server itself and additional information about it may be accessed at the following URL:

<http://ww2010.atmos.uiuc.edu/>

Information on related projects such as HORIZON and CoVis as well as access to their web servers may be found on the UIUC *The Daily Planet* server located at:

<http://www.atmos.uiuc.edu/>

REFERENCES

- Hall, S., et al., 1997: An Online Hybrid Multimedia Environment for K-12 Education in the Atmospheric Sciences. *Proceedings of the 6th Symposium on Education*. Long Beach CA, American Meteorological Society.
- Plutchak, J., et al., 1997: The Weather Visualizer, Java, Habanero, and the Future. *Proceedings of the 13th IIPS Conference*. Long Beach, CA, American Meteorological Society.