Studio 2000 says 'Hello' to the very wide Web World

Lars Bindzus
German National Research Center for Information Technology
Integrated Publication and Information Systems Institute
(GMD-IPSI)
E-mail: bindzus@darmstadt.gmd.de

Steffen Jakob
Vienna University of Technology
Institute of Information Systems
E-mail: S.Jakob@infosys.tuwien.ac.at

September 14, 1995

1 'Welcome to Studio 2000'

You will see this message on top of the Studio 2000 WWW-page. Studio 2000 is a laboratory of the research division DIMSYS (Distributed Multimedia Information Systems) at GMD-IPSI in Darmstadt/Germany, where students mainly belonging to the Technical University of Darmstadt do implementational work.

Multimedia is not a trivial thing. A lot of work and a good hardware equipment are mandatory. That's why the complete institute is equipped very well with a large number of modern workstations; so at Studio 2000. The work is the reason Studio 2000 is permanently occupied by students that are implementing, working on papers or writing their theses.

Multimedia is also a creative thing. New ideas just occur in a funny atmosphere and with new hardware possibilities. In January we got a new SUN Sparc 20 machine with an attached video camera. ... and this camera brought us the stating idea to write a snapshot page as a window to the outside world – to the extent that the world is connected to the Internet.

Since January you can watch us and communicate with us. You only need a connect to the World Wide Web and to Internet Relay Chat (IRC), respectively. Our page is well known now not only in Germany but all around the globe.

2 Concepts

Talking about the World Wide Web is like talking about television or multimedia. You have to see it to get an impression of it. Nevertheless, we will try to describe the whole concept of Studio 2000.

2.1 Studio 2000 Bean – the Logical Structure

The communication concept behind Studio 2000 we call the 'Studio 2000 Bean'. For a graphical overview have a look at figure 1.

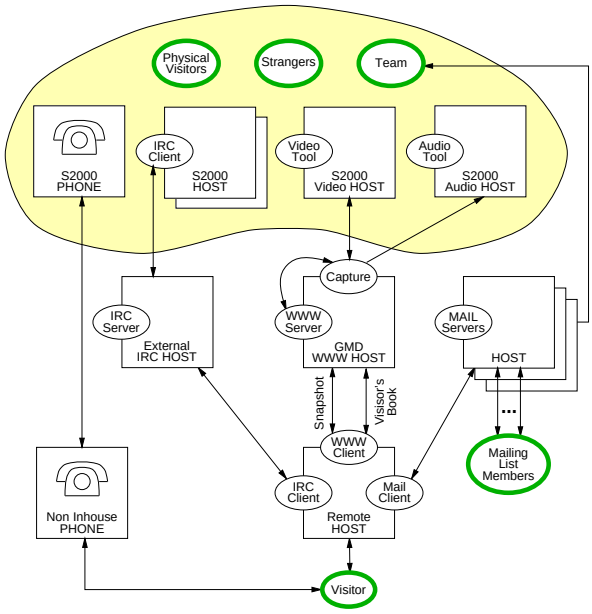

Figure 1: Studio 2000 Bean – A graphical Overview of the Concept

2.1.1 Persons

There are five different groups of persons (shady framed ovals), who are involved in the Studio 2000 world. These groups are described in the following list. Real persons may belong to more more than one group.

Team. The team consists of the following three members, who are also the founders of Studio 2000: Lars *'Schpongo'* Bindzus, Steffen *'Yobes'* Jakob and Thomas *'Checker'* Schecker.

Visitors. Visitors are people who seize the opportunity to participate in Studio 2000. Several ways of communication with persons are described in section 2.1.4.

Physical Visitors. Some persons, who had known us virtually, decided to take a trip to Darmstadt and to become physical visitors. This – of course – does not happen very often.

Strangers. Usually a lot of students or other research workers who are not team members are working at Studio 2000. Nevertheless, they are integrated in our virtual community. Those persons are called strangers.

Mailing List Members. We established a mailing list where everybody can subscribe to. The list gives the friends of Studio 2000 the opportunity to communicate with each other. They are also informed about the ongoing events of Studio 2000. The E-mail address is `s2000@darmstadt.gmd.de`.

2.1.2 Devices

The skeleton of the Studio 2000 world consists of a network of devices (represented by boxes) with different functionalities.

Devices Located at Studio 2000. The heart of Studio 2000 is the live snapshot that will be generated newly each time a World Wide Web-client is calling our page. Therefore we are using a SUN Sparc 20 workstation with an attached video camera (*video host*). Another workstation serves as an *audio host*. We use the audio device of this machine to play sample audio files as a warning signal, when a snapshot is about to be taken. Besides the video and audio host, there are some more workstations in Studio 2000. People who work with these machines are listed in our WWW-page.

The only device which is not connected to the Internet is the *Studio 2000 phone*[1]. Mostly this phone is used for in-house calls. Therefore it might be surprising and even a little bit confusing, when people from all over the world are calling us.

Server Hosts. The server hosts are the most important machines for the transfer of information.

The *WWW-Server* [2] is responsible for the transfer of the WWW-pages and especially of the live snapshot. It runs on a SUN Sparc 10 workstation.

An *IRC-Server* is necessary to build up an online communication via IRC. A lot of such servers exist in the Internet. If you choose a server near at hand (other servers might refuse a connection) you will be able to communicate with all connected IRC-clients.

Access to a *mail server* is indispensable to join the mailing list or to write E-mails to the team.

Remote Hosts. Visitors who want to join the basic Studio 2000 life need to have access to a host connected to the Internet. Another old fashioned way is to call us by phone.

2.1.3 Studio 2000 Laboratory.

The 'real' Studio 2000 laboratory is a room where students of the division DIMSYS[3] are working. In fact this room does not have the shape of a bean as you can see if you have a look at a live snapshot.

2.1.4 Applications and Communication Possibilities

A lot of applications (simple ovals) control the different kinds of communication. An example explains the general ways of communicating with Studio 2000 and the conceptual relationships between the applications, devices and persons:

WWW A visitor starts a *WWW-client* (e.g., Netscape or Lynx). While surfing through the web, he selects a link to the WWW-homepage of Studio 2000. The WWW-server, which runs at the GMD WWW-host, starts the program *Capture*. *Capture* searches for the origin of the request. Depending on the detected Internet-domain or even on the unique host-id *Capture* triggers the *Audiotool* with a specific audio file. All persons in the Studio 2000 room are warned and can prepare themselves for the snapshot. Six seconds later the snapshot will be taken by the *video tool*. *Capture* dynamically creates a WWW-page containing the snapshot, a list of all people who are logged on Studio 2000 hosts, and a lot of other information.

The visitor likes the page and decides to leave a message written as plain text or with included HTML-tags in the visitor's book.

Mail. To subscribe to the mailing list the visitor sends an E-mail to the team. After this he introduces himself to the mailing list members by sending them an E-mail.

IRC. E-mail is not an interactive medium. IRC gives the possibility of online-conferencing to everybody as well as World Wide Web offers the easiest

[1] Phone No: +49 6151 869-900

[2] URL: http://este.darmstadt.gmd.de:5000

[3] Distributed Multimedia Information Systems

way of transferring online-pictures taken by a camera. From time to time the team members join the Internet Relay Chat. Our channel #studio_2000 is not very popular but sometimes funny.

The visitor has announced in his E-mail that he will join the IRC. Therefore all people who are currently present at Studio 2000 and some other people from the mailing list activate their IRC-clients and come together for a chat.

Phone. During chatting, the connection to the IRC-server breaks. The visitor wants to say goodbye to the Studio 2000 team and therefore he calls them by phone. To thank the visitor for his call he gets an entry including a snapshot in the *phone page*.

2.2 Page Layout

The driving idea behind the Studio 2000's page layout can be divided into two different parts: at first one should be able to receive a fancy page even using the simplest graphical WWW-browser. The second motivation was creating a page with a not typically 'computer-like' look.

The first point we fulfilled by using mostly HTML–2 code and the GIF format for all graphical elements. Because the horizontal positioning of the elements is rarely supported in HTML–2, the current layout is a compromise between the features of HTML–3 and the possibility of using old browsers. We figured out a size for the elements positioned in one row where a presentation using an average sized display does not wrap around the particular row.

The second claim resulted from the insight that even if one used all capabilities of the latest graphical design tools, the page still or just therefore might appear in a boring way. The success of perfectly laid-out books is not caused by the latest technology but is a product of the authors' professionalism and experience.

The page layout is relatively simple and our buttons and icons are easy to understand, because our graphical concept ends up in hand-drawn and color-scanned elements which, nevertheless, have a clear signification. Anyway, it is important for hypermedia systems to offer not only graphical, but also textual selection to the visitor. So every button and icon at our page has its corresponding clickable text.

The resulting page is presented in figure 2; it is divided into four screen-sized portions:

First Portion – The Snapshot. The people surfing around the WWW expect pages where they can easily reach the whole functionality of a site without reading too much. That's why the part on top of the page is showing a selection icon-list (where the whole Studio 2000 functionality is reachable from), the snapshot itself and a reference to the last 200 pictures as well as – in our opinion – the best pictures.

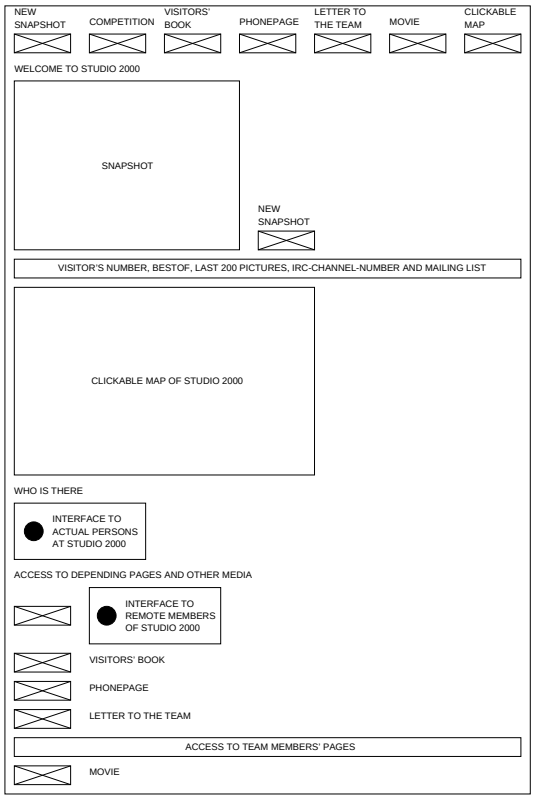

Figure 2: Studio 2000 – The Page Layout

Second Portion – The Studio 2000 Map. A video camera can give you only a limited view of the Studio. So we drew an interactive map, where the visitor can discover all parts of Studio 2000 himself.

Third Portion – List of Actual Workers at Studio 2000. 'Who is actually sitting at a particular workstation?' The visitor will be informed about the actual workers, if they are really working, and can further surf to their home-pages.

Fourth Portion – The Remote Members, Visitor's Book, Phonepage. What is a good WWW-page without a visitor's book? Every visitor can leave a message written as plain text or with included HTML-tags. At this portion of the page we also located references to our remote members, the phonepage and – surely – references to the original team members of Studio 2000.

3 Behind the Scene – Underlying Protocols

This section gives a brief overview of protocols which are the basis for World Wide Web transmis-

sions and Internet Relay Chat data exchange.

3.1 IRC-Protocols.

'The IRC-protocol has been designed over a number of years for use with text based conferencing. *RFC1459* describes the current protocol.'[4]

There are two extensions to this protocol:

DCC. 'DCC allows the user to overcome some limitations of the IRC-server network and to have the ultimate in secure chat connections while still in an IRC oriented protocol.'

CTCP. 'The client–to–client protocol (CTCP) is meant to be used as a way to

1. in general send structured data (such as graphics, voice and different font information) between user clients, and in a more specific case
2. place a query to a user's client and getting an answer.'[5]

3.2 'Hypertext Transfer Protocol'

HTTP is an application-level protocol set up on the TCP/IP of the Internet. It is possible to handle with (multimedia) objects in an efficient and easy way and to construct wide-spread information networks.

The basic idea is a client request and a server response. The client is able to request different actions from the server, e.g., sending, retrieving, updating, deletion of objects in a particular representation as well as executing objects' methods.

For further information about HTTP we recommend the Internet-Draft of HTTP/1.0 that is accessible via the World Wide Web.[6]

4 The Details – Specification and Installation

Now for the boring technical details. This section shall give a brief overview of the specification and implementation idea for a snapshot page like the one of Studio 2000.

The Hardware. We are using a Sparc Station 20 with an attached SUN video camera. Our World Wide Web-server itself is installed on a Sparc Station 10. Actually, the hardware equipment itself is not decisive for a working camera page. You can use any hardware platform if you have an application monitoring a video camera.

Software – The Five Plots. It is not an aim of this article to describe the software installation of a WWW-server. You can find further information about this topic at CERN or NCSA.[7][8] We rather try to explain the specification and installation of the camera software without implementation details.

In the beginning you have to make sure that your WWW-server is allowed to access the particular camera device. Secondly the WWW-server should be able to execute compiled applications as well as those written in PERL or shell-scripts.

The main idea of our specification sounds like this: The client requests our page. Our WWW-server does not send a ready-made HTML-page but is starting a PERL-script. In the script we go through different tasks:

- We update the log files. The first one is a simple file containing the number of the actual requester that will be increased with every request. The second file is a general one, where we just register all information accessible about the incoming request. This information is found in an environment variable generated by the WWW-server at execution time of our script. With the information about the clients, we are able to detect invalid client programs. They will just receive an error message and their requests are naturally not counted.

- We lock the camera while taking the Snapshots. We can guarantee answering as many clients as our WWW-server is able to handle. Furthermore we can guarantee the most actual picture for all requesters! This is achieved by a simple trick:

  It takes about 15 seconds to generate the picture, including conversion (which is the bottleneck as described below). Therefore we write a lock file including the picture filename depending on the visitor number. If another client process requests our page at a later time stamp, it will receive the same picture as written in the file. Then the second process has to wait until the lock file will be deleted, which is done after generation and conversion. This guarantees that all running processes receive a valid picture.

  In case of a system failure the script is enhanced by an automatic timeout of the lock file; that means after 5 minutes the lock file isn't valid any longer and will be overwritten next time the page is accessed.

  During locking another job is done: the picture numbered 200 less than the actual one will be deleted, because we store only the last 200

[4] URL: http://www.cis.ohio-state.edu/htbin/rfc/rfc1459.html
[5] URL: http://www.funet.fi/ irc/docs.html
[6] URL: http://www.w3.org/hypertext/WWW/Protocols/Overview.html
[7] URL: http://www.w3.org/hypertext/www/Daemon/Status.html
[8] URL: http://hoohoo.ncsa.uiuc.edu/docs/Overview.html

snapshots. This should be a reasonable number which, of course, can be enlarged if the requests takes place more often.

- We request an actual picture from the camera. The snapshot itself is done by an application delivered together with the camera and giving pictures in JPEG format as its output. The script is just calling this application.

- We convert our picture to GIF format and process the brightness detection. JPEG would be a proper format if all WWW-clients processed this JPEG. Unfortunately they do not (e.g., old versions of the widespread 'Mosaic' do only know GIF). Annoyingly, the duration of the conversion is a bottleneck for the whole request.

  The conversion task is also the point where our automatic brightness detection is taking place. We use a really simple but ingenious trick: the size of the JPEG-picture. A small size means a highly uniform picture; a high uniformity only appears when it is dark! So we can put out a warning file that 'we are sleeping'!

- We check the people actually working at Studio 2000 and create the HTML-result. After all the HTML-code for the client is sent simply by printing the text to standard output. Most of the text is hard-coded in our script, and only the actual contents are included dynamically (e.g., the picture, time and the visitor's number). At this point the actual information about present people and a dynamic detection of their World Wide Web-pages is also processed.

5 How can I participate?

We hope you are curious to know more about us. You are welcome. First have a look and call the WWW-page:

```
http://este.darmstadt.gmd.de:5000/cgi-bin/capture.pl
```

From within our page you could subscribe to our mailing list. Furthermore, join the IRC-channel:

```
#studio_2000
```

Or just contact us via personal E-mail.

Acknowledgments

We wish to thank Thomas C. Rakow for helpful comments on the manuscript. Moreover we like to thank Regina Klotz for drawing the Studio 2000 Bean. We also thank Thomas Schecker, who is one of the founders of Studio 2000. Last but not least, we want to thank all the nice GMD people and students, who helped us calling the spirit of Studio 2000 into existence.