

Web interfaces

A common feature of many modern GUI applications is that there is a strong desire to deliver the applications via web browsers. MSIE and Navigator have differences in their implementation of (what should be) standard extensions such as Java, and so any web based developments need to be tested on any target platforms.

However - its not really that hard to write CGI/ PHP and/or Java applications that work on all platforms. In this section we will briefly look at CGI, PHP and Java.

7.1 CGI - Common Gateway Interface

CGI is a standard for helping web servers run external programs and return *dynamic* web pages.


For example, a simple dynamic web page might return the current date and time, calculated by running the 'date' program, and formatting the results as a web page. The following script shows the idea:

CODE LISTING	mydate.cgi
<pre>#!/bin/sh cat <<EOM1 Content-type: text/html <HTML><HEAD> <TITLE>Output of data in HTML from CGI script</TITLE> </HEAD><BODY> <H1>Date:</H1> EOM1 date cat <<EOM2 </BODY></HTML> EOM2</pre>	

When this script is placed in the directory `public_cgi` in your home directory on one of the UNIX systems, then you may refer to

```
http://www-cgi.comp.nus.edu.sg:8000/~yourid/mydate.cgi
```

and you will get the following display:



This is of course a trivial example, but shows the fundamental idea of running a script, to get dynamic content in a web page. In this case, the script is not passed any data from the client browser - it just runs the shell `/bin/sh` and the `date` program on the server.

Note that there is no requirement for your CGI program to be a shell script. You may use any suitable scripting language, or compiled programs, and `perl` is very commonly used in this role. The most important thing to remember is that whatever your CGI program does, it should not take too long to process.

7.1.1 CGI environment variables

To pass parameters to CGI programs, environment variables are used. The following **perl** script can display all the environment variables passed to a CGI program:

CODE LISTING	env.cgi
<pre>#!/usr/local/bin/perl print "Content-type: text/html\n\n"; print <<EndOfHTML; <html><head><title>Print Environment</title></head> <body> EndOfHTML foreach \$key (sort(keys %ENV)) { print "\$key = \$ENV{\$key}
\n"; } print "</body></html>";</pre>	

When this script is run, we get something like the following. This may go some way towards explaining to you how some systems know which web browser you are using!

```
DOCUMENT_ROOT = /usr/local/apache/htdocs
GATEWAY_INTERFACE = CGI/1.1
HTTP_ACCEPT = image/gif, image/x-
xbitmap, image/jpeg, image/pjpeg, image/png, */*
HTTP_ACCEPT_CHARSET = iso-8859-1,*,utf-8
HTTP_ACCEPT_ENCODING = gzip
HTTP_ACCEPT_LANGUAGE = en
HTTP_CONNECTION = Keep-Alive
HTTP_HOST = www-cgi.comp.nus.edu.sg:8000
HTTP_REFERER = http://www-cgi.comp.nus.edu.sg:8000/~hugh/
HTTP_USER_AGENT = Mozilla/4.79 [en] (X11; U; Linux 2.2.16 i686)
PATH = /usr/local/bin:/usr/bin:/bin:/usr/local/php/bin
QUERY_STRING =
REMOTE_ADDR = 137.132.90.155
REMOTE_HOST = dhcp-hugh.ddns.comp.nus.edu.sg
REMOTE_PORT = 3343
REQUEST_METHOD = GET
REQUEST_URI = /~hugh/env.cgi
SCRIPT_FILENAME = /home/staff/hugh/public_cgi/env.cgi
SCRIPT_NAME = /~hugh/env.cgi
SERVER_ADDR = 137.132.90.7
SERVER_ADMIN = websp@comp.nus.edu.sg
SERVER_NAME = www-cgi.comp.nus.edu.sg
SERVER_PORT = 8000
SERVER_PROTOCOL = HTTP/1.0
SERVER_SOFTWARE = Apache/1.3.19 (Unix) PHP/4.0.5 mod_perl/1.25
TZ = Singapore
```

The bold strings are the names of the environment variables passed to the CGI program. To the right of the name is the value of that environment variable.

7.1.2 CGI forms

CGI is commonly used for processing simple form-based applications. That is, the client display has a form, and the user keys-in, or selects items in the form, which is then submitted to the CGI program for processing. The principal mechanism for passing small form contents to the CGI program uses environment variables, which are passed to the called CGI program.

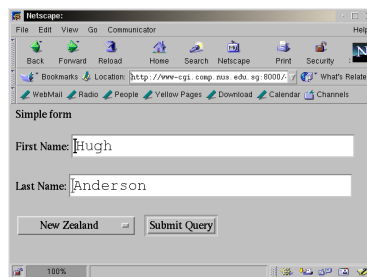
The form contents are found inside an environment variable called **QUERY_STRING**, as a series of **name/value** pairs. This mechanism is known as the **GET** mechanism, and a typical URL would look like this:

```
http://www-cgi.comp.nus.edu.sg:8000/~yourid/myform.cgi?name1=value1&name2=value2
```

An alternative mechanism is the **POST** mechanism, in which the **STDIN** of the CGI program is used to process the form data. A simple example of a form based web page:

CODE LISTING	form.html
<pre><html><head>Simple form</head> <body> <form action="env.cgi" method="GET"> First Name: <input type="text" name="First" size=30><p> Last Name: <input type="text" name="Last" size=30><p> <select name="Home"> <option>Singapore <option>Malaysia <option>Indonesia <option>New Zealand <option>The rest of the world! </select> <input type="submit"> </form> </body></html></pre>	

This produces a page that looks like this:



When the form is submitted, the **QUERY_STRING** looks like this:

```
QUERY_STRING = First=Hugh&Last=Anderson&Home=New+Zealand
```

Within a CGI program, this series of name-value pairs may be used to return a dynamic web page based on this form data. **Perl** is a particularly useful language to use in this context, as it has powerful operators for managing strings, and the **QUERY_STRING** can be **split** quickly into its component parts. There are security issues with unrestricted CGI programs - since they run powerful programs (like perl and csh) with arbitrary parameters, they may be a source of (hacker) intrusion. It is for this reason that CGI usage is restricted here at NUS.

7.2 PHP

PHP is a server-side scripting language that is embedded in your web pages. It looks very like standard HTML scripts, but when a client browser queries a PHP enhanced web page, the server automatically interprets the PHP, and then sends back an ordinary HTML page. There are no enhancements needed for browsers to access PHP web pages.

The two tags `<?php` and `?>` start and end a PHP script, and identify a PHP code segment. The PHP code itself is a reasonably powerful programming language similar to Java, C and Perl, with functions, variables and so on.

PHP stands for PHP - Hypertext Preprocessor, a recursive acronym (like GNU - Gnu's Not UNIX), and is a generally useful HTML/server preprocessor. However it is particularly useful if you wish your web pages to access databases. It is common to pair up PHP with MySQL, but PHP is not limited to one database type. It may be used with any of the commercial databases. For example if you wish to use PHP to access a Microsoft SQL server, you can install the ODBC support in the server machine, and access the server directly.

Here is sample PHP code embedded in a PHP-enhanced web page. It shows PHP connection to a MySQL server, selection of a database and an SQL query:

```
<?php
...
mysql_pconnect("host","user","password")
    or die("Unable to connect to SQL server");
mysql_select_db("dbname")
    or die("Unable to select database");
$numguests = mysql_query("SELECT COUNT(*) FROM guests")
    or die("Select Failed!");
...
?>
```

This may all be integrated with standard HTML and form-based web pages to construct a GUI. PHP suffers less from the security issue than perl or csh CGI scripts do. After writing this sentence I went to see what the latest security issues with PHP were, and discovered that recently there has been a major loophole discovered in PHP POST upload code:

Each of the flaws could allow an attacker to execute arbitrary code on the victim's system.

The exploit appears to have been discovered before any use of it, so assuming you have a relatively recent installation of PHP, you should in general have less security worries than with CGI scripts.

7.3 Java enhanced

In Section 6.5 we saw a very simple hello-world Java applet inserted in a web page. Here is a little Java applet for a Lissajous figure:

CODE LISTING	Lissajous1.java
<pre> /* @(#)Lissajous.java * Original version was written in 0.4 95/04/09 * by Hugh Anderson for HotJava browser. * * Updated by L. Gladney to Java 1.0 JDK on 4/13/97. * * Patrick Chan (chan@scndprsn.Eng.Sun.COM) has suggested that it * would be nice if every point had a different display, so mouse * X motion now controls the ratio of frequencies, and mouse Y motion * controls the amplitude. */ import java.applet.Applet; import java.awt.*; public class Lissajous extends Applet implements Runnable { Thread animate=null; double pi=3.14159265359; int fx=50; int fy=100; int diffx=0; int amp=50,phase=0; // amplitude, phase int delay = 50; // speed set by length // of sleep between refreshes public void init() { resize(200, 200); // resize to fixed width,height } public void paint(Graphics g) { int X,Y,YY=0,lastx=0,lasty=0,temp=0,rev=0; g.drawRect(0, 0, size().width - 1, size().height - 1); // outline if (fy < fx) { // frequency temp = fx; fx = fy; fy = temp; rev = 1; } for (int x = 0 ; x <= 360 ; x += 4) { // loop X = (int) (amp*Math.sin(x*2.0*pi/360.0)); // x pos YY = (x*fy/fx)+phase; Y = (int) (amp*Math.sin(YY*2.0*pi/360.0)); if (x==0) { lastx=X; lasty=Y; } if (rev==1) { g.drawLine(lastx+100,lasty+100,X+100,Y+100); } else { g.drawLine(lasty+100,lastx+100,Y+100,X+100); } lastx=X; lasty=Y; } if (rev==1) { temp=fx; fx = fy; fy = temp; } phase = YY; /* Fix an error ... phase shouldn't increase forever.... */ if (phase < 0) { phase += 360; }; if (phase >= 360) { phase -= 360; }; g.drawString(fx + ":" + fy,10,20); } } </pre>	

Here is the rest of the code...

CODE LISTING	Lissajous2.java
<pre>public void run() { while (true) { repaint(); try { Thread.currentThread().sleep(delay); // delay } catch (Exception e) { }; } } public void start() { if (animate == null) { animate = new Thread(this); animate.start(); } } public void stop() { if (animate != null) { animate.stop(); animate = null; } } public boolean mouseDown(Event e, int x, int y) { Graphics gc; gc = getGraphics(); diffx = fx-x; System.out.println("Got a mouse event at " + x + ", " + y); return true; } public boolean mouseDrag(Event e, int x, int y) { fx = x+diffx; if (fx <= 0) { fx = 1; }; amp = y; return true; } public String getAppletInfo() { return "Lissajous by Hugh Anderson/Larry Gladney "; } public String[][] getParameterInfo() { String [][] info = { {"delay ", "int ", "delay, default=50" } }; return info; } }</pre>	

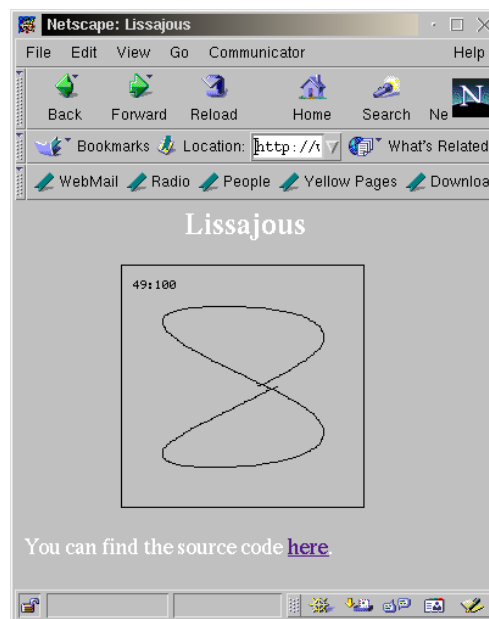
I wrote this code soon after the first Java language was made public, and had forgotten about it until last week, when I went looking for an applet, and found my name on it! This code may be found at

```
http://olddept.physics.upenn.edu/courses/gladney/minicourse/lectures/lecture2.html
```

or locally at

```
http://www.comp.nus.edu.sg/~hugh/Lissajous/Lissajous.html
```

This produces a page that looks like this:



7.4 Summary of topics

In this module, we introduced the following topics:

- Web-based application architectures
 - CGI, PHP and Java
-

Sample questions for this module

1. (Research) Make up a simple CGI form, similar to the one given on page 75, which uses the **POST** method for reading the data, and prints out the three fields.
 2. (Research) Make up a simple PHP processed form, similar to the one given on page 76, and which prints out the three fields.
 3. Examine the `Lissajous.java` code. What is the function of the `diffx` variable?
-

Further study

- <http://php.net>
-

7.5 Sample Assignment 4 - Implementation

Task:

Your task is to implement and document a multi-user simple text editor. By this I mean that you and other people may safely simultaneously edit a single file (perhaps using a shared file system). Note that there is no one correct way of doing this, and feel free to experiment with different approaches.

You may re-use existing simple text editor code found in many introductory Java/Swing tutorials if you wish, but the extensions to make it multi-user must be your own.

Here are brief descriptions of possible approaches to this problem:

1. (Simplest) Each user edits their own copy of the file. When the user saves the file, other users are notified that the original file has changed, and given a choice to load the new changed file, or to just continue.
2. (Medium) Each user edits a part of the file, chosen when they open the file from the remaining editable parts (For example, lines 1-100 for user 1, lines 101-200 for user 2). When the user saves the (part) of the file, other users are notified that the original file has changed, and their programs automatically load the new changed parts.
3. (Very tricky) Each user edits their own copy of the file. As changes are made, all users screens are updated.

Feel free to dream up more advanced shared editing schemes. Note that since communication between the programs may be tricky, I suggest your programs communicate through another shared file.

Deliverables:

- A title page containing your names and matriculation numbers.
- A ten to twenty page document containing
 - A brief summary and justification of the overall strategy used for shared editing - perhaps with a state diagram describing the states of each edit program, with an argument as to why your program is “safe”.
 - An overview of the interface design
 - A user manual for the interface
- A disk containing the code, with a (small) README file.

Visualization

The visualization of data using GUI applications should be distinguished from other computer-graphics concerns. In visualization, we are concerned with *exploration* of data, with its attendant concerns of encoding strategies and so on. In computer-graphics, we may be more concerned with rendering techniques.

Before exploring the implementation of a 3D visualization, we look at some aspects of the context within which the visualization is to be used.

8.1 The use of 3D

A successful visual metaphor has some analog with *real-world* physics. Some studies suggest that a 10-fold improvement in item density can be achieved in using three dimensional displays.

The **etherman** display has proven effective in observation of networks with 50 or less nodes, but becomes cluttered and unusable with more nodes on-screen. By extending the display into the third dimension, it immediately becomes clearer. Our familiarity with spatial location allows us to understand that objects further away will be smaller, and this reduces the visual clutter. However if a *far away* object increases in size, we immediately notice, and can rotate the display to observe more closely. We notice even if the *far away* object is still smaller (in screen *real-estate* terms) than closer objects. This human cognitive behaviour becomes apparent as soon as sufficient visual cues have been given to persuade the observer that the display is in three dimensions.

For example: Figure 8.1 shows the output of an original program¹ to display tasks active on a UNIX machine. The size of the spheres indicate the amount of memory used by each process, the colour represents the owner. This display can be rotated and used to examine activity in a way unattainable using standard system process viewing tools. The display has over 100 visible nodes, but it is still easy to identify and investigate individual nodes.

¹<http://opo.usp.ac.fj/~hugh/Public/Viz/ThesisWork/processes1>





Figure 8.1: Display of tasks in a multi-tasking environment.

8.2 OpenGL

OpenGL was originally the SGI in-house graphics system, but now is the most widely accepted graphics standard, with chip, API and OS support for all platforms. It is possible to code directly using the OpenGL API, but more normal to use a toolkit which encapsulates some abstraction, built on top of OpenGL calls.

Open GL is standard on all UNIXes and all versions of Windows since Win95. The API supports functions for rendering, buffering, anti-aliasing, shading, colouring, texture-mapping, a display list, Z-buffering and so on. To give the flavour of raw OpenGL programming, here is a small application:



CODE LISTING	teapot.c
<pre> #include <GL/glut.h> void Teapot (long grid) { /* ... code to construct drawlist of teapot here. */ } static void Init (void) { glEnable (GL_DEPTH_TEST); glLightModelfv (GL_LIGHT_MODEL_LOCAL_VIEWER, local_view); /* Lighting model, materials... */ } static void SpecialKey (int key, int x, int y) { switch (key) { case GLUT_KEY_UP: rotX -= 20.0; glutPostRedisplay (); break; /* Move in other directions */ } } static void Draw (void) { glClear (GL_COLOR_BUFFER_BIT GL_DEPTH_BUFFER_BIT); glPushMatrix (); /* ... translations ... */ glCallList (teaList); glPopMatrix (); glutSwapBuffers (); } int main (int argc, char **argv) { glutInit (&argc, argv); type = GLUT_RGB GLUT_DEPTH; type = (doubleBuffer) ? GLUT_DOUBLE : GLUT_SINGLE; glutInitDisplayMode (type); glutInitWindowSize (300, 300); glutCreateWindow ("TeaPot"); Init (); glutReshapeFunc (Reshape); glutKeyboardFunc (Key); glutSpecialFunc (SpecialKey); glutDisplayFunc (Draw); glutMainLoop (); } </pre>	

8.3 Java 3D, VTK - toolkits for 3D

These systems are 3D OO toolkits embedded in Java and C++ respectively. The Java 3D application programming interface (API) provides a set of object-oriented interfaces that support a simple, high-level programming model.

The Visualization ToolKit (VTK) is an open source OO software system for 3D consisting of a C++ class library, and several interface layers for Tcl/Tk, Java, and Python. VTK has a wide variety of visualization and graphical functions, and has been installed and tested on both UNIX and Windows.

8.4 Case study - network traffic application

A user-requirement specification for a network traffic application may begin with something like:

This visualization is to assist network managers in planning and monitoring their networks. It allows interactive exploration of network datalink traffic, and is intended for use both for visualization of immediate-mode (real-time) data, and for visualization of historical data. (The visualization is the same in each case, except that time only travels forward in the immediate mode.)

The visualization will help answer questions such as the following:

- *Which segments carry the most traffic?*
 - *Which sections of the network are down?*
 - *At what times, and where do traffic bottlenecks occur?*
 - *What is the line utilization for different lines at different times?*
 - *What types of traffic are used most?*
 - *Would routing or switching be effective here?*
-

For this *network* traffic application, the following elements are represented:

- Background: - to convince the viewer that the display is *three dimensional*...
- Nodes: - a computer, a network device...
- Traffic: - the amount of traffic flow...
- Protocol: - the *type* of traffic...
- Errors: - errors in traffic could be further traffic *protocols*...
- Trends: - for changes over time...
- Association: - for network insights...

8.4.1 Node representation

In our chosen context, the nodes represent computers or network components such as hubs, routers, bridges or switches. In locational or representational displays we may want to differentiate between the type of node, but in the more *abstract* displays, there may be no need to do this.

In Figure 8.2 we see a range of possible options for more concrete representations of nodes.

The computer represented in Figure 8.2(a) has about 2000 flat triangular surfaces (some of them hidden). If we were visualizing a campus with (say) 500 computers using this representation,




Figure 8.2: Concrete node representations.

Machine	Rendering speed	Computer (a)	Computer (b)	Computer (c)
Graphics Workstation	485,000 Δ /sec	0.485 frames/sec	11.5 frames/sec	69 frames/sec
PC1	30,000 Δ /sec	0.03 frames/sec	0.71 frames/sec	4.3 frames/sec
PC2	11,000 Δ /sec	0.011 frames/sec	0.26 frames/sec	1.6 frames/sec

Table 8.1: Workstation redraw speed.

then our rastering engine has to recalculate the positions and shading of 1,000,000 polygons each time it redraws the screen. This will happen even if the item is so *far away* that it only takes up a single pixel on the screen.

A typical modern hardware rastering engine can calculate 485,000 shaded Δ /sec (triangles per second), and hence our screen refresh rate would be about half a frame per second, giving a *jerky* look. By contrast, the computer shown in Figure 8.2(c) has only 14 flat triangular surfaces, giving a frame rate in excess of 70 frames per second.

Standard PCs often come with graphics cards that support pixel movement on screen, but their overall performance in shaded Δ /sec is normally considerably below 485,000 polygons per second. Table 8.1 gives the resultant frame rates for displaying onscreen 500 of the node representations in Figure 8.2.

It is clear from this table that if we wish our visualizations to be viewed on a range of platforms, we must choose our node representations carefully to minimize rendering time.

Some representation methods for three dimensional objects allow different *levels of detail*. In the VRML specification, a single object may be represented in different ways depending on how much screen *real estate* it uses up. If the object is near you, it could be represented in detail, but if it is a long way away, the representation could be as simple as a coloured square.

The following VRML code represents a cone in two ways using an LOD (Level Of Detail) node. If the distance from the user to the object is smaller than the first range value specified, then the first version is drawn. If the distance is greater than the last range specified, the last version is drawn.

```
#VRML V2.0 utf8
LOD {
  range [20]
  level [
    #full detail 16 sided cone
    Shape{
      appearance Appearance { material Material { diffuseColor 1.0 1.0 1.0}}
      geometry Extrusion{
        crossSection [ -1 0, 0 0, -1 -2 -1 0]
        spine [1 0 0 , 0.866 0 0.5, 0.5 0 0.866, 0 0 1 , -0.5 0 0.866, -
0.866 0 0.5,
                -1 0 0, -0.866 0 -0.5, -0.5 0 -0.866, 0 0 -1 ,0.5 0 -0.866,
                0.866 0 -0.5, 1 0 0]
      }
    }
    #low detail 4 sided cone, actually a pyramid
    Shape{
      appearance Appearance { material Material { diffuseColor 1.0 1.0 1.0}}
      geometry Extrusion{
        crossSection [ -1 0, 0 0, -1 -2 -1 0]
        spine [1 0 0 , 0 0 1, -1 0 0, 0 0 -1 , 1 0 0]
      }
    }
  ]
}
```

8.4.2 Traffic and protocol representation

A simple immediate way to represent traffic between two nodes is to just draw a line between them. The nature of network communication on a typical Local Area Network (LAN) is such that the resultant lattice is likely to be relatively sparse.

For example: at the datalink layer, on average, a workstation at any one time may only be communicating with six or seven other datalink addresses - two broadcast addresses, (say) two file servers, a WINS or DNS server. and a proxy. So - rather than having a lattice with $\frac{n^2-n}{2}$ interconnections, we have $k(n-1)$ interconnections, where k is some small integer. Even so, a lattice with 500 nodes may have 3,000 interconnections and may look jumbled.

A line indicates source and destination, but not the *amount* of traffic. Three systems for this purpose have been examined:

1. Colour coding (black through red to white for maximum traffic),
2. Line width, and
3. The length of partial lines, as discussed in Eick's papers.

Using a linear increase in the line width appears most effective, although it does increase clutter. It also leaves the colour information free for use in some other encoding. The linear scale needs a sensible maximum, and experimentation has shown that a maximum width should be equivalent to the size of the node.

A simple line or cylinder also does not tell which way the traffic is flowing. We evaluated the following cues by modeling them in *geomview*, a geometrical modelling package.




Figure 8.3: Partial length representation of bi-directional traffic.

1. Separate arrows
2. Partial lengths

In Figure 8.3 we see the traffic between two computers, the size of the cylinder between the machines indicating the total amount of network traffic, and the two colours indicating the relative amounts of traffic going each way. The nodes and cylinders themselves are coloured according to the dominant protocol type.

8.4.3 Trend representation


Trends are sometimes difficult to find in large sets of data such as found in our application. Once an examination of a visualization has indicated that a trend may be possible, it is normally easy to frame the questions needed to verify the trend.

- “*It looks like HTTP usage on these segments is increasing...*”
(\Rightarrow Plot HTTP usage for the segment machines versus time).
- “*It looks like HTTP usage is increasing when FTP usage is decreasing...*”
(\Rightarrow Plot HTTP and (1-FTP) versus time).

Graphing continues to be the pre-eminent way of representing trends and the role of visualization is to assist in finding the trends.

The four-dimensional visualization methods outlined and demonstrated by Olaf Holt and Nils McCarthy in NDdemo (the fourth dimension being explicit time) could perhaps be used in trend analysis, but the visualization is a little hard to use.

A final method is to attempt to encode previous visualizations *on-top-of* the current one (but perhaps semi-transparent) - the idea here is one of *visual* echoes. In only some circumstances can this be successful. There are two options:

1. Echoes are fixed on the screen, and we can move the visualization away from them, leaving a trail like this: 
In the worst case though, we have just ended up using one of our three display dimensions for “*time*”.




Figure 8.4: Locational view

2. Semi-transparent echoes are co-located with the visualization. In this case, we can only show some of the history. We can show a reducing item, but not an increasing one.

8.4.4 Display

In Figure 8.4, we see an early locational view showing the fixed components of the visualization, and modeled using **geomview**. It shows nodes for the computers, floor plans for the buildings, and a transparent roof. The display uses the normal 3D navigation tools for adjustment. From the display it is easy to identify the location of machines, and the display should be efficient enough to support display and manipulation of the entire network (with 500 machines as a suitable goal), and - yes - the computers are floating in mid air. (Since we are concerned with efficiency we choose the simplest understandable visualization, and tables just become extra polygons to draw).

The visualizing tool supports rotation and translation of the display, so that the observer can easily focus on regions of interest. Suitable systems are found in the CosmoPlayer VRML viewer, and in **geomview**. Note that this visualization is not dependant on the navigation or implementation method.

Cables and network infrastructure are not marked on the display, but the display does support an aggregation-by-rule construct. This aggregation can be used to associate machines all on the same segment, or all used by the same department.




Figure 8.5: Aggregation nodes

The components are chosen to be the best minimum complexity representation consistent with fast updates. The frame update speed for the most complex display is better than 2 frames per second.

Each node or aggregation in the overview display is clickable to turn it off or on. If a node is turned off, its traffic no longer is displayed (either directly or as part of some aggregated traffic). If an aggregation is turned off, any existing traffic displays to its nodes are removed. This facility is used to allow fast reduction of visual clutter.

Aggregation nodes also have an aggregation *switch*, which allows them to combine all traffic for subsidiary nodes. When this switch is on, lines connect the aggregation to its subsidiary nodes.

The aggregation node floats above its associated nodes. In Figure 8.5, we see two aggregation nodes, with the one on the right aggregating traffic to and from all its subsidiary nodes. All traffic to or from these nodes is displayed going to the aggregation node. The other nodes are not aggregated, and display traffic directly.

Each node or aggregation in the overview display is clickable to identify specific information about that node. This information does not replace the display, but appears in a separate window. Initially this information may just be textual information such as the name of the node along with traffic totals, but eventually, it is expected that the drill-down display will be the metaphor display specified elsewhere, showing only the selected node in 3D, along with any associated nodes.

Operating System	Web Browser	VRML 2.0 Plugin
IRIX	Navigator 3.01S	CosmoPlayer 1.0.2b3 or later
	Communicator 4.04	CosmoPlayer 1.0.2b3 or later
	Communicator 4.07	CosmoPlayer 2.1 beta
Macintosh	Communicator 4.04	CosmoPlayer 2.1 or later
WIN32	Navigator 3.01	CosmoPlayer 1.0 beta 3 or Intervista WorldView 2.0 or later
	Communicator 4.04	CosmoPlayer 1.0 beta 3 or Intervista WorldView 2.0 or later
	MSIE 3.0	CosmoPlayer 1.0 beta 3 or Intervista WorldView 2.0 or later
	MSIE 4.0	CosmoPlayer 1.0 beta 3 or Intervista WorldView 2.0 or later

Table 8.2: Systems which support the EAI.

8.5 3D VRML visualization implementation

The VRML visualizer is a relatively small Java program which must be loaded as an applet along with a VRML view of the network. A small web page is created, and may be used to view the visualization using a web browser such as Netscape along with the CosmoPlayer VRML plugin.

Unfortunately, not all combinations of web browser and VRML plugin work correctly with the EAI, but the systems in Table 8.2 are known to work. These systems were current in 1999. This year (2002) all the systems I tried at NUS appeared to work fine.

Load the default web page in the directory, and the visualization should be visible. To finish using the visualizer, you must exit the browser entirely. If not, the Java applet keeps communicating with the **collector**.

In Figure 8.6, we see an active VRML display within a browser. The computer nearby is generating a lot of traffic. In the distance we can see other nodes, and the roof and floors.

8.5.1 3DVNT VRML software

3DVNT includes software to create a default HTML web page for the VRML visualization. The current default web page is like this:

```
<html><head> <title>Sample 3DVNT Page</title> </head>
<center><H1>Sample 3DVNT Page </H1></center>
<center> <embed src="root.wrl" height="600" width="700"> </center>
<center> <applet code="View1.class" width="100" height="10" mayscript>
<PARAM name="segment" value="MACS"> <PARAM name="port" value="9876">
<PARAM name="host" value="opo.usp.ac.fj"> </applet> </center>
OK?
</html>
```




Figure 8.6: 3DVNT view within Netscape browser.

The *root.wrl* file which forms the basis of the VRML visualization is of the following format:

```

PROTO CLUSTER [] { ... }      # Cluster definition
PROTO KEYBOARD [] { ... }    # Keyboard definition
PROTO SCREEN [] { ... }     # Screen definition
PROTO GLOBE [] { ... }      # Traffic sphere definition
# Some setting up declarations
Background { skyColor .4 .66 1 }
NavigationInfo { type [ "EXAMINE", "ANY" ] speed 400 }
Viewpoint { position 0 400 0 orientation 0 1 0 4 description "Camera 1" }
# Lines, floors and roofs
DEF LINES Transform { ... }
DEF FLOORS Transform { ... }
DEF ROOFS Transform { ... }
# and then the nodes
DEF node1 Transform { ... }
DEF node2 Transform { ... }
# ... and so on ...

```

Each node is of the following form:

```

DEF node1 Transform {
  translation 4350 150 4365
  rotation 0 1 0 4.71238
  children [
    KEYBOARD {}
    SCREEN {}
    DEF node1box Transform {
      children [
        Shape {
          appearance Appearance { material DEF node1boxcolor Material { diffuseColor 0.8 0.8 0.8 } }
          geometry Box { size 50 50 50 }
        } ] ]
    DEF node1sphere Transform {
      scale 1 1 1
      children [
        Shape {
          appearance GLOBE {}
          geometry Sphere { radius 1 }
        } ] ] ] ]
} ] ] ] ]

```

The Java visualizer software maintains a link to a remote data collector, and uses the EAI to modify the images in the VRML view.

Mar 05, 99 11:51	View1.java	Page 1/3
------------------	-------------------	----------

```

// using the VRML External Interface.

import java.applet.*;
import java.awt.*;
import java.util.*;
import vrml.external.field.*;
import vrml.external.exception.*;
import vrml.external.Node;
import vrml.external.Browser;
import java.io.*;
import java.net.*;

public class View1 extends Applet {
// public static final int DEFAULT_PORT = 9877;
    Browser browser;
    Socket s = null;
    DataInputStream in = null;
    String line;

    public void init() {
        System.out.println("Test.init(...)");
    }
    void SocketStart () throws java.io.IOException {
        String port = this.getParameter("port");
        int p = Integer.parseInt(port);
        try {
            String host = getCodeBase().getHost();
            System.out.println("Request came from: " + host);
            s = new Socket(host, p);
        }
        catch (UnknownHostException e) {
            System.out.println("No socket: " + e);
        }
    }
    public void start() {
        int count=0;
        Node node2sphere=null;
        Node appear=null;
        EventInSFVec3f[] scalein=new EventInSFVec3f[100] ;
        EventInSFColor[] appears=new EventInSFColor[100] ;
        float[] val = new float[3];
        int[] lastval = new int[100];
        int n;
        String id,v1;

        while (count != 100) {
            scalein[count] = null;
            appears[count] = null;
            lastval[count] = 0;
            count=count+1;
        }
        try {
            SocketStart();
        }
        catch (java.io.IOException e) {
            System.out.println("No socket: " + e);
        }

        System.out.println("Test.start(...)");
        browser = (Browser) vrml.external.Browser.getBrowser(this);
        System.out.println("Got the browser: " + browser);

        count = 0;
        try {
            in = new DataInputStream(s.getInputStream());

```

Mar 05, 99 11:51	View1.java	Page 2/3
<pre> tEventIn("scale"); tValue()); tIn("set_diffuseColor"); ntValue()); </pre>	<pre> while(true) { line = in.readLine(); if (line == null) { System.out.println("Server closed connection."); break; } if (line.regionMatches(0,"n",0,1)) { n = line.indexOf(32,2); id = line.substring(2,n); System.out.println(">>>" + id + "<<<"); vl = line.substring(n+1); System.out.println("+++ " + vl + " ---"); Integer a = Integer.valueOf(id); Integer b = Integer.valueOf(vl); if (scalein[a.intValue()] == null) { try { node2sphere = browser.getNode("node" + id + "sphere"); System.out.println("Got the sphere node: " + node2sphere); } catch (InvalidNodeException e) { System.out.println("PROBLEMS! node2sphere: " + e); } try { scalein[a.intValue()] = (EventInSFVec3f) node2sphere.ge System.out.println("Got the sphere scale node: " + appears[a.in } catch (InvalidNodeException e) { System.out.println("PROBLEMS! (scalein): " + e); } try { appear = browser.getNode("node" + id + "boxcolor"); System.out.println("Got the Boxcolor node: " + appear); } catch (InvalidNodeException e) { System.out.println("PROBLEMS! appearance: " + e); } try { appears[a.intValue()] = (EventInSFColor) appear.getEven System.out.println("Got the Boxcolor color node: " + appears[a.i } catch (InvalidNodeException e) { System.out.println("PROBLEMS! appearance color: " + e); } } if (b.intValue() == -1) { val[0] = (float)1.0; val[1] = (float)1.0; val[2] = (float)1.0; } else { val[0] = (float)(b.intValue()*20)+1; val[1] = (float)(b.intValue()*20)+1; val[2] = (float)(b.intValue()*20)+1; } scalein[a.intValue()].setValue(val); if ((b.intValue() == 0) != (lastval[a.intValue()] == 0)) { if (b.intValue() == 0) { val[0] = (float)0.8; val[1] = (float)0.8; val[2] = (float)0.8; appears[a.intValue()].setValue(val); } } } } </pre>	<p>2/3 Thursday August 26, 1999</p>


```
Mar 05, 99 11:51 View1.java Page 3/3
    } else {
        if (b.intValue()==-1) {
            val[0] = (float)0.1;
            val[1] = (float)0.1;
            val[2] = (float)0.1;
            appears[a.intValue()].setValue(val);
        } else {
            val[0] = (float)0.0;
            val[1] = (float)1.0;
            val[2] = (float)0.0;
            appears[a.intValue()].setValue(val);
        }
    }
    lastval[ a.intValue()]=b.intValue();
}
//      } System.out.println(line);
}
} catch (IOException e) { System.out.println("Reader: " + e); }
}

public Browser getBrowser() {
    return browser;
}
}
```

Thursday August 26, 1999 3/3

8.6 Summary of topics

In this module, we introduced the following topics:

- Visualization versus computer-graphics
 - OpenGL
 - (Briefly) Java3D, VTK
 - VRML/Java/EAI
-

Sample questions for this module

1. Find a minimal VRML file which constructs a solid cube.
 2. Find minimal OpenGL display-list code to draw a cube.
 3. Examine the Java code given for using the EAI to modify the VRML. How exactly does the code get a reference to a VRML node?
 4. Examine the Java code given for using the EAI to modify the VRML. How exactly does the code modify a VRML node?
-

Further study

- sunsite for Java3D
 - The EAI specification
-