Geotraceman: A Visual Traceroute

Scott Pleitner & Darren Brown
Department of Computer Science
Curtin University of Technology
P.O. Box U1987, Perth 6001
Australia
email: scott@cs.curtin.edu.au
June 19, 1995

Abstract

Geotraceman is a network monitoring tool that aids in understanding the position of hosts throughout the Internet. Geotraceman achieves this by drawing trace lines on a map of the world. These trace lines represent the path Internet Protocol (IP) packets take to get from one host to another.

Geotraceman also provides the delay statistics for each trace. Geotraceman uses the X Windowing System to provide its' output in a simple to comprehend and use interface.

The map can display trace information in either a flat (cylindrical projection) map or on a three dimensional globe. Additional map functionality includes zooming and rotation.

1.0 Introduction

The Internet is a vast community of internetworked networks. As such, it can be difficult to understand, in a geographical sense, where each host resides. It is quite a simple task to discover (in logical terms) the section of the Internet (subnet) where a host resides, but this tells us nothing of its geographical location. Knowledge of the location of a subnet still may prove useless in locating an individual host.

Geotraceman is a member of the Netman suite of tools developed by the Department of Computing Science at Curtin University of Technology. The concept behind this on going project is "*effective network diagnostics that are within the financial reach of network managers within small organisations*" [11].

2.0 Background

The Internet is a community of computer networks linked together. The fact that no single governing body is required indicates the monumentous co-operation achieved by the Internet community.

There is, however, a single body that administers the Internet.

"The Internet is loosely administered by the Network Information Center (NIC) at the Stanford Research Institute (SRI), where documents and standards are archived, IP address are allocated and network and domain names are registered" [10].

The lack of a governing body and *Internet Police* highlights the understanding, trust and co-operation exhibited by the Internet community.

2.1 Protocols Used in the Internet

For any network to operate successfully, all connected hosts must be able to communicate. This is achieved through the use of communications protocols common to all hosts. The Internet uses the TCP/IP (Transmission Control Protocol/Internet Protocol) suite of protocols. This protocol suite includes numerous protocols for transmitting data and control information. Geotraceman operates through the use of a combination of these data and control protocols.

The protocols used directly by Geotraceman are IP (Internet Protocol), UDP (User Datagram Protocol) and ICMP (Internet Control Message Protocol).

2.1.1 Internet Protocol

The Internet Protocol is the main protocol used throughout the Internet. It acts as the outer shell for almost all other protocols that flow through the Internet. As stated in Comer [3]:
"The fundamental Internet service consists of a packet delivery system. Technically, the service is defined as an unreliable, best-effort, connectionless packet delivery system... The service is called unreliable because delivery is not guaranteed. The packet may be lost, duplicated, delayed, or delivered out of order, but the service will not detect such conditions, nor will it inform the sender or receiver. The service is called connectionless because each packet is treated independently from all others. A sequence of packets sent from one machine to another may travel over different paths, or some may be lost while others are delivered. Finally, the service is said to use best-effort delivery because the Internet software does not discard packets capriciously; unreliability arises only when resources are exhausted or underlying networks fail."

Packets that are unable to be delivered to their destination host do not remain in the Internet forever. Each IP packet has a time-to-live (ttl) field which determines the longevity of the packet. When a packet arrives at a router the value of this field is decremented then checked. If it is zero, the packet is removed and an error message is sent to the source of the packet. If the time-to-live is not zero, the field is theoretically decremented by the time the datagram resides within the router and the time taken to travel to the router. (This does not occur in practice, hence the time-to-live field degenerates into a counter indicating the number of sites touched by the packet - a hop count.) The packet is then routed appropriately.

2.1.2 Internet Control Message Protocol

ICMP is a protocol which generates control packets used within the Internet. This protocol "*allows gate-*

ways to send error or control messages to other gateways or hosts" [3]. ICMP was originally designed for router to router control, but can also be used to transmit such control messages between hosts.

ICMP packets are used to retrieve trace data from the intermediate and destination sites. The data returned in these packets provides a method of calculating the delay statistics and the route taken.

Two ICMP packets of particular interest to Geotraceman are those containing the *Destination Port Unreachable* and the *Time Exceeded for Datagram* messages. The *Destination Port Unreachable* message indicates the destination host has been reached (see section 2.2 Traceroute to clarify this apparent contradiction) and the *Time Exceeded* message indicates an intermediate site between the source and destination (once again, see section 2.2 Traceroute). Statistics are calculated as the return ICMP packets are received, hence a real-time comprehension of the delays that can occur in a packet switched, wide-area, network can be realised. Even though ICMP packets are control packets, they are still subject to the same connectionless, unreliable delivery from which all Internet traffic suffers.

2.1.3 User Datagram Protocol

Comer [3] succinctly summarises UDP:
"The User Datagram Protocol provides unreliable, connectionless delivery services using IP to transport messages between machines. It adds the ability to distinguish among multiple destination within a given host computer."

UDP is the primary mechanism for application programs to transmit datagrams to other application programs. This is achieved through the use of protocol ports. UDP provides no error detection, error correction or flow control. It is assumed that the application program will handle this.

The main advantage of UDP is its' small, eight byte, header. This makes for extremely efficient data transfer when compared to the other major transport protocol used in the Internet (Transmission Control Protocol or TCP) which has a header of at least 24 bytes (depending upon options). TCP provides reliable end-to-end communications with flow control and error detection. These features are not required by Geotraceman.

2.2 Traceroute

Traceroute [7] performs the task of tracing a route between hosts. This is achieved by attempting to force an ICMP *Time Exceeded for Datagram* response from each router on the path to the destination host. The ICMP *Time Exceeded for Datagram* message is produced when a router examines an IP packet with a time-to-live field of zero.

Traceroute uses a UDP packet as a probe datagram. This datagram is sent to port 33434 which is, hopefully, unused on the destination host. The first packet traceroute sends is a 38 byte UDP probe packet contained within an IP packet with a time-to-live field of value one. When the first router receives this packet, it decrements the time-to-live. Since this field now contains zero, an ICMP *Time Exceeded for Datagram* message is sent back to the source. The second packet sent contains a time-to-live field of two. The first router receives this packet, decrements the time-to-live field and, since it is non-zero, forwards it. The second router receives the packet, decrements the time-to-live field and sends an ICMP *Time Exceeded for Datagram* message back to the source host as this field now contains zero. This process continues until the destination host receives the probe datagram. As it is unlikely a process will be using port 33434, an ICMP *Destination Port Unreachable* message should be generated by the destination host and sent back to the source of the packet. This is the end of the trace.

The current implementation of traceroute does suffer from some problems. These problems include port 33434 being in use on the destination host and the route taken by outbound packets changing during a trace. If the required port is in use on the destination host, an ICMP *Destination Port Unreachable* message will not be generated. The process that is using this port will receive the traceroute probe datagram and treat it as data. In essence, this prevents the trace from completing. Traceroute provides no guarantee that all outbound packets will take the same route. Routing decisions can occur during a trace, dynamically modifying the route. The current traceroute has no method of detecting, nor solving this problem.

2.2.1 A New Traceroute

A new version of traceroute has been defined by Malkin in RFC 1393 [8]. This version is more efficient than the current implementation, requiring only $n+1$ packets rather than $2n$ (where n is the number of hops between the source and destination). This new version utilises an ICMP message created specifically for tracing routes.

Rather than repeatedly sending packets with different time-to-live values, the new traceroute sends a single packet to the destination. At each site along the path, the new ICMP reply packet is generated. This addresses one of the major problems with the existing traceroute: all outbound packets need not take the same route. The information generated at each hop is free to take whichever route desired back to the source as it is the information within each packet which contains the required trace statistics. As a UDP probe datagram is not required in Malkin's version of traceroute, there is no port required on the destination host, hence eliminating the port in use problem with the current traceroute.

2.3 Maps

The map data used by Geotraceman was extracted from the CIA World Map Databank [4]. The CIA World Map Databank is a vector based data set which uses longitude and latitude to store points. This data was modified using a customised thinning algorithm designed specifically to reduce the size of the data while maintaining as much detail as possible.

Rather than treating the entire map as a single region, it is divided into smaller, more manageable regions. The scale of the coordinates is then reduced as all references to points can be made relative to their region, rather than the entire map. A vector can be addressed by selecting the region it lies within and then simply using offsets within this region to indicate the vertices. This allows for efficient zooming as, where possible, clipping is performed on region boundaries. Entire regions of the map can be removed from consideration, increasing processing speed dramatically.

The original map data was over 20 megabytes in size. This was reduced down to the 6 megabyte dataset Geotraceman currently uses. The reduction algorithm uses a threshold distance between points. Any points within the threshold distance are ignored. For a more complete description of the map reduction, please refer to Brown [2].

Two options were considered for accessing the data:- minimising disk access or minimising memory usage. These two options are mutually exclusive as minimising memory usage requires the data to be loaded from disk when needed and discarded from memory when not, and minimising disk access requires the data to be read only once, but stored in memory for the duration of the program. Due to the size of the dataset, the minimising memory usage option was chosen.

The map data for Geotraceman exists at three levels, each of differing detail. As required, each file is loaded from disk and operated on. When no longer required, the memory used to hold the data is released.

2.4 Interface

Geotraceman uses the X Windowing System [9] as its interface. X is a networked graphical user interface which compliments the requirements of Geotraceman. X is the most common graphical user interface found on UNIX workstations which is the target machine of Geotraceman.

"X is a distributed, intelligent, device independent, operating system independent, hardware independent windowing system"
(Personal communications - Andrew Marriott, Senior Graphics Lecturer, Curtin University of Technology)

2.4.1 Graphics Library

Rather than rely on *raw* X calls, an external graphics library was used. This simplified the coding process, but decreased the speed of display as an extra level of abstraction was added. The graphics library used was VOGL (Very Ordinary Graphics Library) [5]. This library was a close emulation of the Silicon Graphics GL library [11]. The majority of features used in GL are available in VOGL. The functions which receive the most use from Geotraceman are the line drawing functions and the double buffering functions (for efficient drawing).

3.0 Quad Trees

As with most large projects, execution speed is important to Geotraceman. After profiling, the map displaying was found to be the most time consuming aspect of Geotraceman. Due to this, various methods were considered in an attempt to increase the speed of display. Using a quad tree to store the map data was an option.

Quad trees were not implemented in Geotraceman due to the increase in the size of the dataset. The speed improvements in actual display times would have been noticeable, yet the time taken to load the much larger dataset would possibly have offset this. Some testing with quad trees indicated an increase in the size of the dataset by approximately 2.2 times (allowing for all of the pointers used to reference branches in the tree).

3.1 Quad Tree Basics

Quad trees are a recursively decomposable data structure. They provide multiple resolutions, dependent upon the level of detail required. This complements the zooming features of Geotraceman as less data can be displayed at a lower zoom and more at a higher zoom. Zooming speeds could also be increased as the zoom area could be clipped efficiently by the quad tree itself.

Quad trees are based on a *divide and conquer* problem solving techniques. Each branch of the tree can be subdivided, traversed and *conquered* independently of the branches above. This can be accomplished effectively through the use of recursion.

A quad tree consists of two elements, leaf nodes and non-leaf nodes. Leaf nodes contains actual data and are normally found towards the bottom of the tree. Non-leaf nodes contain four branches to other nodes (either leaf or non-leaf). Non-leaf nodes often hold a summary of the tree below for quick reference.

3.2 Quad Trees in Geotraceman

After consideration of the time taken and the difficulty encountered in producing the initial data set, it was decided not to alter the map data to produce a quad tree. The solution to the original problem of display speed was instead solved by maintaining the current data set and manipulating it more efficiently.

The main concern about display speed arose during implementation of the three dimensional projection. The introduction of clipping produced a 50% reduction in the amount of data to process, hence the time taken to display the globe was greatly reduced from initial expectations, further reducing the need for quad trees.

4.0 Traces

A trace is made up of connected line segments between two end points:- the trace source and destination. This line indicates the path an IP packet *could* take to travel from the source to the destination. This is not a guaranteed path as IP packets are routable and routes can change without notice.

4.1 Two Dimensional Traces

A trace on the flat map consists of a series of straight line segments connected at sites, signified by filled circles. The names of the sites *actually* touched are written around this circle. How these names are resolved is discussed later in section 7.0 Geographical Location Resolution. The trace line segments represent the shortest distance between two adjacent sites based on the current rotation of the map. A line segment that is expected to travel east may actually travel west if travelling east involves breaking the edge of the visible map. This can be seen in Figure 1: Flat Map with a Trace. The physical path taken from Australia to America is eastwards, passing over Hawaii, rather than the westward path over Africa as shown.

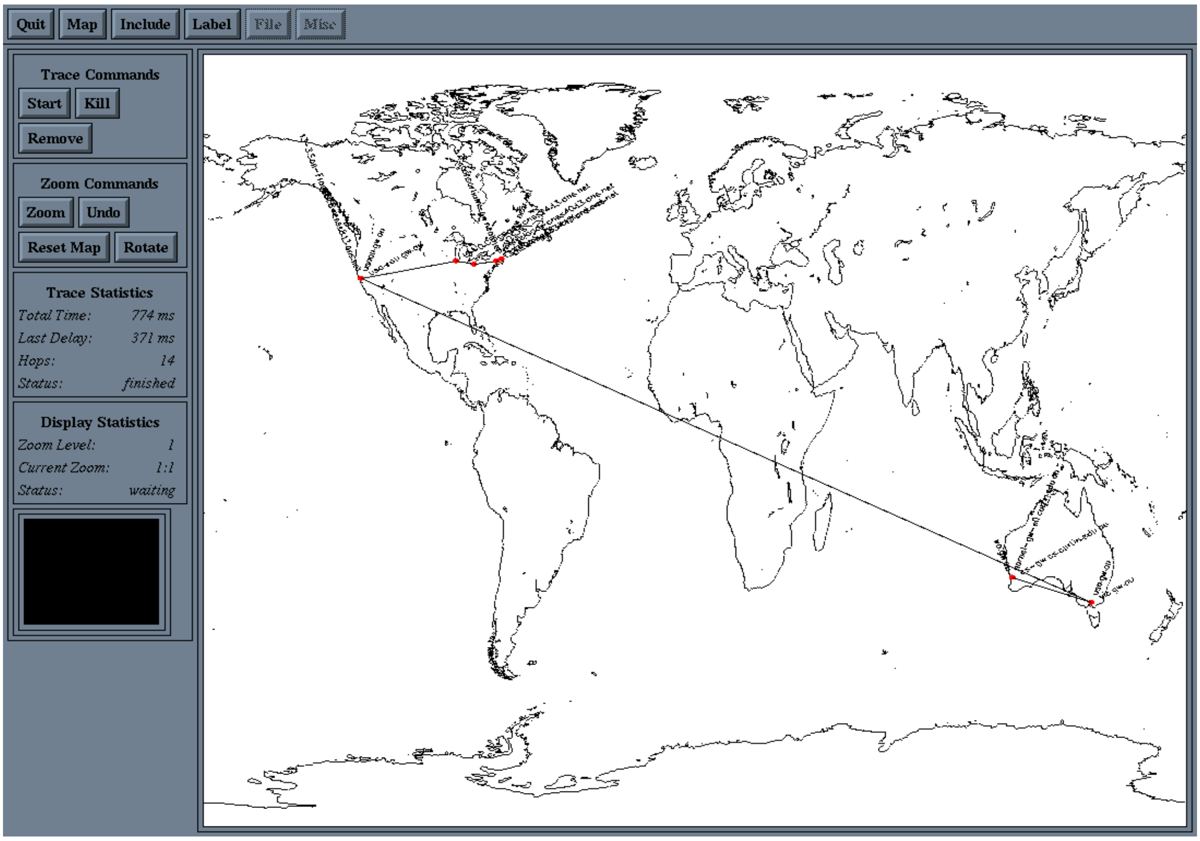

Figure 1: Flat Map with a Trace

The trace itself is performed by spawning a *traceroute* [7] process and sending the received data back to the main Geotraceman process via a non-blocking, unnamed pipe. This data is then parsed for the site name and the delay statistics. The site name is resolved and displayed, the trace line segment drawn and the delay statistics calculated/updated and displayed.

4.2 Trace Statistics

The trace statistics returned by *traceroute* represent the delay each packet takes to get from the source host to each intermediate site/destination host. This is represented in Geotraceman as the *"Last Delay"* statistic. The other delay statistic displayed by Geotraceman is the *"Total Delay"*. This represents the largest amount of time for each hop, summed to produce the maximum time taken by the total trace. Another piece of information available with the trace statistics is the "*Trace Status"*. This informs the user if a trace is currently underway. The trace status can be *tracing*, *finished* or *waiting*. The state is toggled by events such as a trace starting, a trace finishing or a trace being removed.

5.0 Map Functionality

The Geotraceman map attempts to provide the *natural* functionality expected by the user. The map attempts to imitate an atlas as closely as possible. It aims to provide the same level of functionality as an atlas with the added ability of allowing the user to define their own map views (within commonsensical boundaries). This functionality includes such features as zooming, rotation and map overlays.

Geotraceman provides map statistics relating to the current zoom. These statistics include the "*Zoom Level*", the "*Zoom Ratio*" and the "*Map Status*". The "*Map Status*" is used to indicate when a zoom is being performed as detailed zooms can take a few seconds to complete. The "*Zoom Level*" is an integer indicating the detail level. As the zoom level increases, numerous extra features are included such as more detailed maps and extra cities (see sections 5.4 Map Inclusion and 5.5 City Inclusion below). The "*Zoom Ratio*" is, as its name suggests, the ratio of the current zoom.

5.1 Zooming

Geotraceman has the ability to zoom into any area of the map at almost any zoom level. As the map data for Geotraceman is available at different levels, the exact display is dependent upon the zoom level. At a higher zoom level, more detail is displayed.

Zooming is implemented via a bounding box. The box is drawn on the map canvas by the user. Once drawn to the required scale, the new zoom is calculated, the map statistics are updated and the zoom is drawn.

5.1.1 Bounding Box

The zoom function utilises a bounding box for specifying the zoom extents. This box is maintained as a square, simplifying the process of producing a zoom which maintains aspect. The square aspect of the bounding box is calculated as the bounding box is

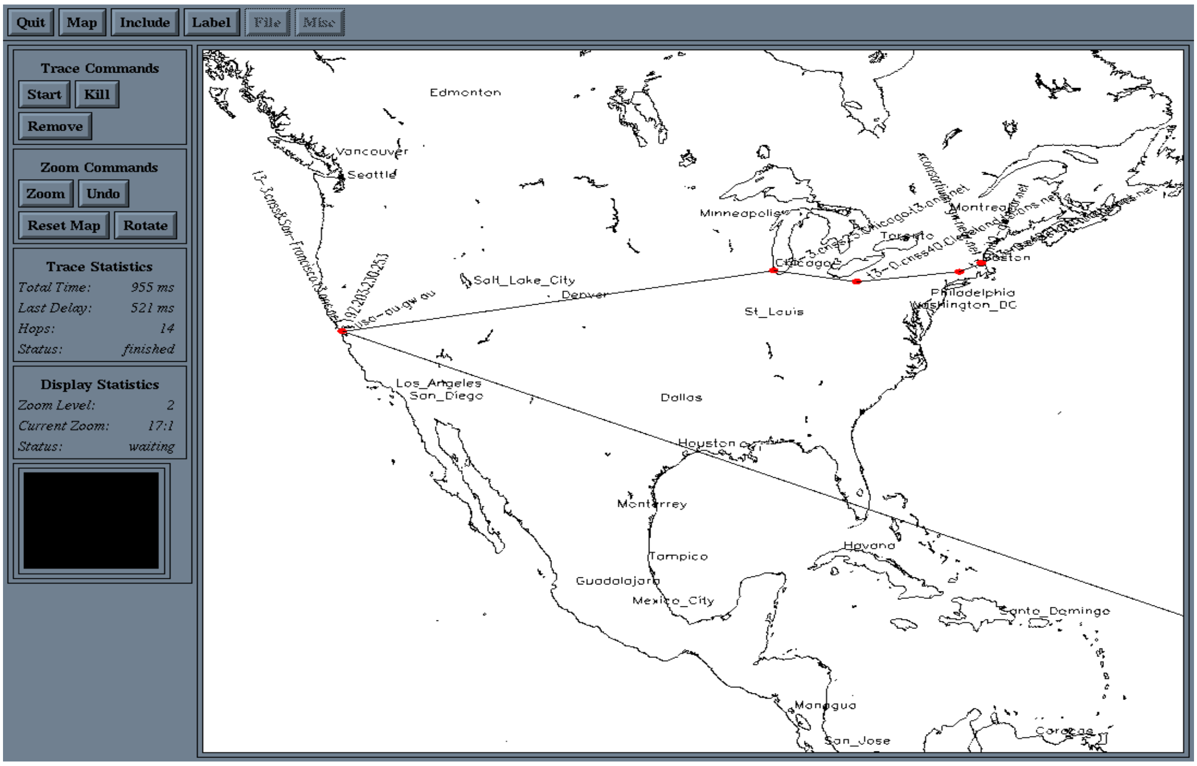

Figure 2: Zoomed Two Dimensional Trace

dragged across the map canvas. As expected, the bounding box has the ability to mirror itself across both the x and y axis. The square aspect is maintained even through mirroring. Figure 2: Zoomed Two Dimensional Trace shows a close up zoom of the USA including the same trace from Figure 1.

5.2 Unlimited Undo

An unlimited undo feature has been incorporated into Geotraceman to allow the user to reverse any modifications made to the map. The undo is implemented via a stack to allow the user access to the previous changes in the reverse order to which they occurred. This follows the *natural* model outlined above.

The data stored for the unlimited undo is all of the necessary map information required to recreate any combinations of zooms, rotations and map inclusions.

An added feature to the undo is the ability to reset the map to its original state, undoing all modifications made. This provides the user with a method of resetting the map without having to undo all of the changes made.

The design of the undo was aimed mainly at the situation where a zoom is accidentally performed. Originally, Geotraceman only allowed a zoom to be performed by selecting the appropriate button, then drawing the bounding box on the canvas. This was modified to allow zooming at any time, simply by drawing the bounding box on the canvas. This increased the probability of accidental zooms. The undo option was seen as a method of allowing the user to restore their previous state without having to recreate all of their zooms, rotations and inclusions.

5.3 Rotation

Rotation involves the apparent movement of the map in the direction specified. This is achieved by modifying the location of the bounding edges of the currently displayed map, then redrawing.

Currently, rotation is only available in the east and west directions. The ability to rotate in a north and south direction would be a welcome addition to the *natural* functionality of Geotraceman. It should be possible to rotate in these directions when the rotation would not extend the map beyond the poles (there can be no wrap-around in the north-south direction as there is in the east-west direction, unless a globe is considered).

Geotraceman performs its rotation in jumps 1/10 of the map extents currently displayed, e.g. if the map is at full zoom, 1/10 of 360° is 36°, hence rotation is performed in 36° jumps. If the map is zoomed slightly to show only 200° of the world, rotation is performed in 20° jumps.

5.4 Map Inclusion

The map inclusion features of Geotraceman allows the overlaying of thematic maps (maps with a single theme) over the coastlines map. This provides greater detail and can aid in maintaining a point of reference for the user.

Geotraceman offers two standard overlays: rivers and political boundaries. These two overlays are drawn in different colours directly over the coastlines map when required. They are in exactly the same format as the coastline map, hence are displayed using exactly the same routines. This allows the overlays to be rotated

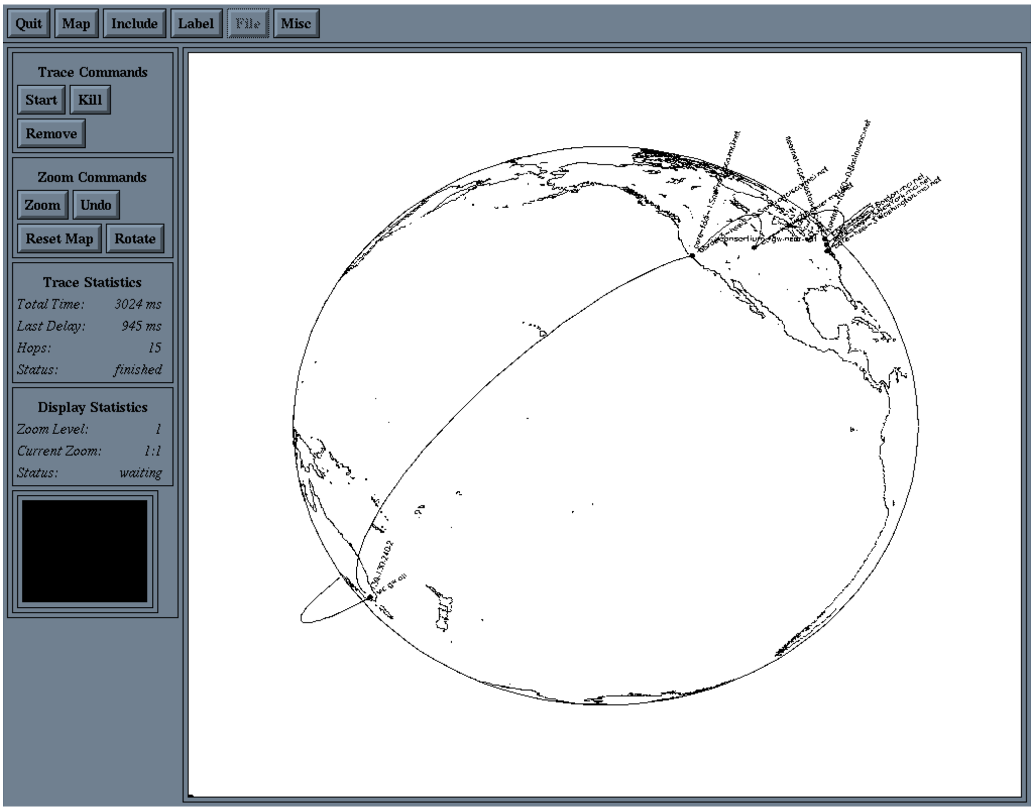

Figure 3: Three Dimensional Projection

and zoomed in unison with the coastline map without having to employ specialised functions.

5.5 City Inclusion

The ability to include cities is also offered by Geotraceman. The exact number of cities displayed is dependent upon the number of cities found in the cities files, the location of these cities and the current zoom level. As the level of zoom increases, a more detailed cities file is read in. The location of the city and those around it determines if it will be displayed. If a city is already displayed in the immediate vicinity, the newly found city will not be displayed. This prevents city names from overwriting each other.

The basic algorithm used to determine if a city should be displayed is outlined below:

```
Boolean clip_city
BEGIN
for all cities already displayed do
   if ((abs(testcity.x-BUFFER)≤oldcity.x)
    AND(abs(testcity.y-BUFFER)≤oldcity.y))
      return TRUE
   else
      return FALSE
end_for
END
```

```
display_city
BEGIN
if (zoom_level   1)
   while(NOT end_of_file(cities_file_1))
      read city from cities_file_1
      if (clip_city = TRUE)
         draw city
   end_while

if (zoom_level   2)
   while (NOT end_of_file(cities_file_2))
      read city from cities_file_2
      if (clip_city = TRUE)
         draw city
   end_while

if (zoom_level   3)
   while (NOT end_of_file(cities_file_3))
      read city from cities_file_3
      if (clip_city = TRUE)
         draw city
   end_while
END
```

From the above, it can be seen for any level of zoom, greater than one, the city files for all the previous levels are also displayed. This allows each city to be entered only once, rather than once in each file. This complicates the above algorithm, but reduces the chance of having the same city entered into each city file at slightly different locations.

The buffer specified in the above algorithm is the amount of space to be left between adjacent cities. Currently, this is hard-coded. In future releases, it will be a user configurable resource.

Duplicate cities can still be displayed despite the best attempts of the above algorithm. If the city in question is specified at two locations further apart than the buffer zone, both be displayed. Clipping can only be performed on the location field as many city names are duplicated across countries, states and even within states.

Figure 2: Zoomed Two Dimensional Trace illustrates the above algorithm for a zoom level of two.

6.0 Three Dimensional Projection

The three dimensional projection offers no extra functionality over its' flat map counterpart. It is simply another viewing method. It allows the user to visualise a trace over a globe rather than the flat map.

Release 1.0 of Geotraceman does not support the displaying of trace lines in three dimensions. They are still displayed as a straight line rather than a curved arc as expected in a three dimensional model. This has since been addressed, but not released.

Production of the three dimensional projection was performed in two main steps: clipping and drawing.

6.1 Clipping

Clipping the data before processing increasing the speed of display enormously. The majority of time consumed by the three dimensional model was not actually displaying the data, but preparing the data for display.

A 50% clipping was achieved through some simple, inductive logic. At any one time, only half of the globe is visible, hence only half of the data needs to be considered. This clipping produced an enormous time saving.

There is, however, a limitation to the clipping. Both end points of all vectors are passed through the clipping routine and if one end point passes the test and the other does not, the entire vector is not drawn. A solution to this would be to perform an intersection test to find the exact location of the border crossing. This would produce perfect results, but would decrease the speed of the clipping. At the time of writing it was felt that the visual improvement was not worth the trade-off in speed.

6.2 Drawing

Many attempts were made to produce a three dimensional projection from the two dimensional data. The algorithm which proved most successful is outlined below.

```
for each point displayable do
   alpha=point.x*mapwidth*2*π
   beta=point.y*mapwidth/2*π/2
   point.x=radius*cos(beta)*sin(alpha)
      +mapwidth/2
   point.y=radius*sin(beta)+mapwidth/2
end_for
```

Through some efficient coding, the above algorithm can be made to execute extremely fast. The number of calls to the maths library can be reduced to only two and the entire algorithm can then be written in terms of multiplications and additions. The effects of this algorithm can be seen in Figure 3: Three Dimensional Projection.

The above algorithm performs its operation in two steps. The first is a conversion from pixel based longitudes and latitudes into degree based coordinates. These are then projected onto a three dimensional plane in preparation to be drawn onto a two dimensional display plane.

The ability to produce multiple projections has been considered and will most likely be included in a future release of Geotraceman. Producing the map canvas as a separate widget (see section 8.0 Future Plans) would allow the simple addition of any desired projection.

During the development of the initial three dimensional projection, a failed attempt produced a projection called Mollweide's Projection. While this was a failure at producing the required projection, it was felt that this still conveyed some useful information. The inclusion of Mollweide's Projection will occur in the next release of Geotraceman, as will the ability to display traces over it.

7.0 Geographical Location Resolution

The process of finding a sites geographical location is not simple. There is no central, nor distributed, store of such information. One method used in Geotraceman is to store this data locally in cache files. The constant modification occurring in the Internet quickly makes any static data obsolete. This data is no different. Another approach taken involves each site maintaining location data of their site locally. Other hosts could then query this information via the Domain Naming System (DNS).

Neither of these solutions produces perfect results. They are both subject to the accuracy of the input data and the ability to find this data. If a hosts' location can not be resolved, it is placed at the same position as the previous host. This is a far from an ideal solution.

7.1 Cache Files

The location of sites in the Internet is maintained by Geotraceman in a series of three cache files. The three cache files are the backbones, the networks and the local caches. The backbones cache holds the mapping of all known backbone sites to geographical locations. The networks file maintains a list of known networks (represented by their netmask) and their geographical location. The local cache file contains local hosts whose geographical position is different from the rest of its' network.
The backbones cache file is used for major routing sites in the Internet. Such sites could include the major (now defunct) NSFNet, AARNet and JANet sites and network switching points. The networks cache file contains entries for networks such as CUT-NET (the Curtin University network) with a netmask of 134.7 and CUT-NET Kalgoorlie (a geographically different campus) with a netmask of 134.7.60.

The search order for the cache files is backbones followed by local, followed by networks. If the required site is not found in any of these files, it is assumed to be at the same location as the previous host. This can lead to some large inaccuracies if the cache files are not well maintained, but no method has been discovered to alleviate this.

7.2 Domain Naming System

"The Domain Name System is a distributed database. This allows local control of the segments of the overall database, yet data in each segment are available across the entire network through a client-server scheme. Robustness and adequate performance are achieved through replication and caching" [1].

DNS consists of two main parts: the name server and the resolver. The name server answers requests generated by the resolver. In the context of Geotraceman, *named* (the name server daemon) is the name server and Geotraceman itself (via the resolver libraries) is the resolver. A query is generated by Geotraceman and sent to the name server. If possible, the name server extracts the required data from its' database and replies to the resolver. If the data is not available, the resolver is informed accordingly.

As the DNS currently does not define a resource record which specifies geographical positioning, a new definition was required. To this end, the GPOS (geographical position) RR was defined in RFC 1712 [6]. The data contained in this RR is longitude, latitude and altitude.

A query for the GPOS information can only succeed if the local systems administrator has entered the data. If the query is unsuccessful, Geotraceman falls back on the cache file resolution outlined above. This still does not solve the problem of an unresolvable host.

8.0 Future Plans

Geotraceman is far from a complete tool. There are many modifications and extensions that should be added before it is considered finished. A preliminary release of Geotraceman has already taken place. This release only supported the flat map and cache files.

The aim of the preliminary release was to receive feedback from the Internet community about the usefulness of the tool and, hopefully, updated cache files. Geotraceman has been adopted by the Internet community fairly well with nearly 1000 ftp's of the source and binaries in the first three months since its release.

Some possible extensions to Geotraceman are outlined below:

- Postscript Dump - The ability to dump the contents of the map canvas to a postscript file for later printing would be a valuable addition. The majority of Netman tools supports postscript dumps. (This has since been completed, but not included in any release to date).
- Multiple Traces - Rather than only being able to display a single trace, showing multiple traces could aid in understanding routing and the "backbones" of the Internet.
- Trace Loading and Saving - Retrograde monitoring can be used for comparative and educational purposes. Traces that were performed months apart can be compared and students can be given access to the tool without requiring "root" permission to start a trace.
- Site List - Maintaining a list of sites touched along with some brief statistics about each site could prove useful in debugging a wide area network problem. This could be used to provide more detail than the summary trace statistics currently present.
- SNMP Monitoring and Management - Adding SNMP (Simple Network Management Protocol) monitoring and management features to Geotraceman would increase its functionality enormously. Geotraceman could now be used to correct problems that it detects. When a routing problem is detected, an SNMP query could exactly locate the problem, then an SNMP set could be used to correct it.
- Map Widget - By encapsulating the map displaying functions of Geotraceman into a widget, it allows for easy incorporation into other tools. This also subdivides the task of improving the performance of Geotraceman. Plans are underway to include the Geotraceman map into other Netman tools such as Interman and other new tools, currently under development. (The Map Widget is currently being developed, with completion date around November 1995).

9.0 Conclusions

Geotraceman has proven to be a useful tool in the areas of education and wide area network error detection. Geotraceman can be used as an educational tool to provide a better understanding to the geographical extent of the Internet. It also provides an insight into the operation of a packet switched wide area network. Geotraceman can, and has, also be used to detect routing errors in the Internet. It detects such situations as router failure, bad routes and cyclic routes. Geotraceman does not report these errors, but rather leaves it to the network manager to infer the error from the data presented. At this stage, Geotraceman can not correct such errors, but with the inclusion of SNMP management, it should be possible.

10.0 References

1. Albitz, P. & C. Liu, DNS and BIND, O'Reilly and Associates, Inc., California, 1992
2. Brown, D., Geographical Trace Route v1.0, Technical Report, Department of Computer Science, Curtin University of Technology, 1993
3. Comer, D.E., Internetworking with TCP/IP, Volume 1, Principles, Protocols and Architectures, Prentice Hall, Englewood Cliffs, New Jersey, 1991
4. Dellinger, J., README file supplied with the CIA World Map Databank, Stanford University, 1989
5. Echidna, E., VOGL - Very Ordinary Graphics Library, Department of Engineering Computing Resources, University of Melbourne, 1991
6. Farrell, C., M. Schulze, S. Pleitner & D. Baldoni, DNS Encoding of Geographical Location, RFC 1712, Network Information Centre, SRI International, November, 1994
7. Jacobson, V., The *traceroute* manual page, 1989
8. Malkin, G., Traceroute using an IP option, RFC 1393, Network Information Centre, SRI International, January, 1993
9. MIT X Consortium, X11 Windowing System Release 5, Massachusetts Institute of Technology, Massachusetts, 1991
10. Nemeth, E., G. Snyder & S. Seebass, UNIX System Administration Handbook, Prentice Hall, Englewood Cliffs, New Jersey, 1989
11. Schulze, M., G. Benko & C. Farrell, Homebrew Network Monitoring: A Prelude to Network Management, Technical Report, Department of Computer Science, Curtin University of Technology, 1993
12. Silicon Graphics Inc., Graphics Library Reference Manual, C Edition, Documentation Version 4.0, Silicon Graphics Inc., Mountain View, California, 1990