Network Mapping and Usage Determination

by F.P. Senekal and J.S. Vorster

Keywords

Network Mapping, Host Discovery, Topology Determination, Usage Determination, Port Scanning, Service Detection, OS Detection, IP Geolocation, Visualization, Geographical Information Systems

Abstract

A large computer network such as the Internet contains millions of computers, services and users, interconnected in a complicated and ever changing web. This article provides an introduction to network mapping and usage determination – the study of the physical connectivity of a computer network and how it is being utilised.

The article discusses the different processes involved in gathering information about computer networks, such as host discovery, port scanning, service detection and operating system detection. Based on this information, topology determination techniques can be applied to infer network structure from the information. Techniques to visualise the information are discussed. IP geolocation (the ability to associate a geographical coordinate with a node in a network) and its applicability to visualising network information on geographical information systems are described. The research is illustrated by means of software tools that have been developed for this purpose.

1. Introduction

Today, there exist large corporate, academic and governmental computer networks containing literally millions of computers distributed over a large geographical area. The ability to monitor and control these networks starts with an understanding of how these networks are structured and function. The objective of this article is to introduce the reader to the techniques and activities involved in determining the structure and function of such large computer networks and to provide some intuition as to how such information may be visualised.

First, some of the equipment that is typically found in a computer network is introduced. The problem of how to uniquely identify the equipment is addressed. Next a number of techniques that are used to discover equipment in a network, such as pinging, Media Access Control address resolution and alias resolution are introduced. Topology determination – techniques for discovering how the different equipment is interconnected – is described, as well as some of the potential pitfalls of topology determination techniques.

Apart from the structural layout of a computer network, it is also of interest to determine the functional elements of a network. Usage determination is discussed, introducing techniques such as port scanning, service detection and operating system detection.

Once the structural and/or functional components of a network have been determined, the information may need to be visualised in order for a human to build understanding and insight about the network. Some abstract layout algorithms are introduced, as well as IP geolocation, a technique by which equipment in a network is mapped to geographical coordinates.

2. Computer Network Equipment

There exist various types of equipment in a computer network. The most obvious example is the computers that people use every day to carry out their daily activities. These computers are connected to computer networks in order to benefit from the information, processing or workflow abilities that are available on other equipment in the computer network. These abilities are usually provided by dedicated server computers. The server computers expose services that are available to the computers that connect to it (often called the client computers or hosts). In between the client and server computers, there exist various other types of computer equipment that are used

to distribute and route the information between the various computers. Such equipment can typically interconnect multiple pieces of other equipment through communication links. Routers and switches are intelligent devices that are able to analyse the data that arrives at them and forward it on an appropriate communication link. They are typically used to connect various networks together to form a Wide Area Network (WAN), but are also found in configurations where they are used to interconnect the computers in a Local Area Network (LAN). Hubs and bridges are less intelligent devices that forward the data it receives on all the communication links connected to it. They are typically found in LANs. Apart from the equipment that is used for the distribution of information, other equipment is used for the purposes of security. An example is a hardware firewall. Such a firewall typically has two network connections. When it receives data on one communication link, it will only forward it on the other communication link if the data satisfies certain criteria that are configured on the firewall.

For the purposes of this discussion, the different pieces of equipment in a network will be known as nodes, irrespective of the type of equipment. Nodes are physically connected by communication links. The set of nodes with their interconnections define what is known as a graph in graph theory. The reader is referred to the book by Epp (Epp, 1995) for an introduction to graph theory.

A central question underlying this work is how a node is uniquely identified. Nodes exhibit various properties, such as a name and one or more associated Internet Protocol (IP) and Medium Access Control (MAC) addresses (see Stallings, 2000 or Kurose, 2001). Properties such as these may be used to uniquely identify a node. These properties may however change over time, and may thus not be adequate to capture the identity of a node; however, there currently exists no better method. In practise, the authors found that the use of the IP address works best in uniquely specifying a node.

Another point to be made is that any network mapping algorithm takes a fixed duration to execute. For large networks, the network topology and usage may change during the time that the network mapping algorithms are executing. Thus, at the completion of the execution of the algorithms, the result is merely a reasonable approximation to the true state of the network at that time. For the purposes of the discussion, the assumption is made that the network mapping algorithms are executing or that scans are conducted from computers that are connected to a network.

3. Host Discovery

Host discovery is the process of discovering which nodes are located in a network. It is also referred to as scanning a network. A number of techniques that are used in host discovery are discussed.

A. Internet Control Message Protocol Echo Command

The Internet Control Message Protocol (ICMP) (RFC792, Postel, 1981) is used to communicate network layer information over a network. It is often used in identifying problems in the communication environment. It provides for the usage of echo and echo reply messages (collectively know as pinging) to be exchanged between two nodes in order to establish whether communication between two nodes is possible.

A source node creates an ICMP echo message and sends it out. The message contains the IP address of a destination (target) node. The message is routed to the destination node that in turn interprets it and creates an ICMP echo reply message. The reply message is routed back to the source node. When such a message is received at the originator of the message, it is interpreted that communication is possible between the two nodes. It can also be deduced that a node with the target IP address existed and was operational at the time the scan was conducted (of course, it may have been possible for a node in between the source and destination nodes to have generated a positive response, which would certainly be a rare and unexpected event).

If a positive reply is not received, it can be for a variety of reasons. For instance, there may be a node or communication link between the source and destination that is not operational, the destination node may be down, the destination node may actively prevent a reply on the message or the message may be actively dropped by one of the nodes in the path. The latter case happens for instance where Internet Service Providers (ISPs) drop ICMP echo messages on their networks due to the additional bandwidth and cost implications.

Pinging is the most commonly used technique for host discovery and is versatile in the sense that it can easily be used to check a range of IP addresses.

B. Media Access Control Address Resolution

With each network interface card (NIC) installed on a node in a computer network is an associated Medium Access Control (MAC) address. A MAC address consists of six bytes. The first three bytes indicate the manufacturer of the card and the last three bytes is a unique serial number assigned to the NIC by the manufacturer. Each NIC in the world is supposed to have a unique MAC address (although there are NICs where the MAC address is configurable).

MAC addresses are used in the routing of messages and can be determined by the Address Resolution Protocol (ARP) (RFC 826, Plummer, 1982). ARP functions in the same way as ICMP echo messages, except that the actual MAC address is returned as part of the reply. Nodes for which a MAC address is returned are thus deduced to be operational.

An advantage of ARP is that reply messages are almost always generated for operational machines, since the process is inherit to how data is distributed across a network. The drawback however is that the ARP only resolves IP addresses for nodes on the same LAN, which severely reduces its applicability to large networks.

C. Alias resolution

For human use, the numerical identity of IP and MAC addresses can be awkward to work with. Humans as better at remembering names than numbers. As such, nodes in a network often have one or more names or aliases associated with them. Examples include hosts aliasing, where names are assigned to nodes in a network, mail server aliasing, where names are assigned to mail servers in a network, and web server aliasing, where a name is associated with a website that is hosted on a web server.

In order to translate these names into and IP address, Domain Name System (DNS) (RFC1034 and RFC1035, Mockapetris, 1987) servers exist. These servers are distributed geographically and can be queried through DNS query and reply messages.

A DNS query is almost always guaranteed to provide an IP address for a valid name in the server database, since the process similar to MAC address resolution in that it is fundamental to the way most networks operate. Alias resolution is limited in the sense that the alias needs to be known beforehand. As such it is often used for footholding purposes, for example, given that the name of a corporation is known, likely names for the associated website names and mail servers may be guessed. These names are resolved to determine IP addresses, which can be analysed to determine a range of IP addresses that can be queried through other techniques.

D. Other techniques

Port scanning (see the section on usage determination) is another technique that can be used to discover nodes on a network. If a connection can be made to a port on a specific IP address, it implies that a node with that IP address was operational at the time that the scan was made.

Packet sniffing is a passive technique where the messages that flow across a network are listened to. It implies access to a network and a device that is capable of intercepting messages. Through analysis of the messages, equipment on the network and the connections they make to other equipment can be inferred.

E. Considerations

The question of which nodes could possibly exist in a network need to be addressed. IP addresses are assigned to organisations by the Regional Internet Registries (RIRs). For instance, the Council for Scientific and Industrial Research (CSIR) have been assigned the block of IP addresses from 146.64.0.0 to 146.64.255.255. This implies a total of 65536 IP addresses that could be used within the CSIR's computer networks. Only a fraction of these IP addresses may be assigned to nodes in a network at any given time, however; it does place a limit on the number and range of nodes that need to be investigated.

Another consideration is the completeness of the of the network information gathered, i.e. to what extent all the nodes in a computer network have been discovered. Techniques such as pinging and port scanning could be directed at all the nodes in a network, which could in theory reveal all the nodes in the network. On the other hand, techniques such as alias resolution and packet sniffing would only reveal some of the nodes in a network. Of course, any technique would only reveal the nodes that are operational, powered and accepting of the scanning protocols. Some nodes, such as hubs, are invisible to the scanning protocols. This is due to the unintelligent nature of these devices, in that they do not interpret the network traffic that passes through it. More advanced techniques are required to detect such devices in a network.

Another issue is how readily a scan conducted against a network can be detected. In theory, all active techniques (techniques that introduce network traffic) can be detected. Pinging and port scanning are active techniques that can easily be detected. Alias resolution and MAC address resolution are active techniques, but less prone to detection due to the fact that they are inherit in the way networks operate. Packet sniffing is a passive technique and is least prone to detection.

4. Topology Determination

Network topology determination is the process of determining the complete structure of a computer network, i.e. the complete set of nodes and communication links between nodes. The host discovery process described earlier can be used to determine the nodes in a network, but not the links in the network. Multiple good articles exist on topology determination (Huffaker, 2002, Hyun, 2003, Lowekamp, 2001, Siamwalla, 1999).

A communication link serves as a channel to carry a transmission signal and cannot be 'queried' like a node. Therefore, the existence of a communication link must be inferred from the nodes between which the link exists.

This can be achieved by means of successive applications of the ICMP echo message, commonly known as a traceroute. A special field can be set in an individual ICMP echo message that specifies the maximum number of hops before the message 'expires'. A hop is equivalent to transmission of data over a single communication link, thus a specification of *n* hops means that the message will expire after it has traversed *n* communication links. Every node in the communication chain decreases the value of the amount of hops traversed in the message. When this value is decreased to zero, the node at which the final decrement occurs generates a return

message to the originator of the message. The return message contains the IP address of the node at which the message terminates.

A traceroute works by sending out an ICMP echo message with a maximum hop value of one, waiting for a reply, sending out a message with maximum hop value of two, waiting for the reply, etc. The generation of echo messages terminate either when a successful reply is generated from the target node, or when one of the echo reply message is not received. Such a reply may not be received for any one of the reasons discussed in the section on host discovery by means of ICMP echo messages.

A single traceroute thus provides information about which nodes are possibly connected together in a chain. It may be inferred that nodes that are differ by a single hop in a traceroute are connected via a communication link. This is generally true, but there may be exceptions as the example in figure 1 illustrates. In the example, a traceroute is done from node A to H. The first ICMP echo message is limited with a maximum hop count of one, the result being B. The second message is limited by a hop count of two and routed from B on its way to H across the C-D nodes. At C, the maximum hops expire and a message is sent back to A containing the terminating node C. A third ICMP message is sent. This time however, B may choose to route the information to H across the E-F nodes, for instance due to changing traffic conditions. The result will be that F is returned as the node with a hop count of three away from A. A similar procedure would result in G being a hop count of four and H being a hop count of five away from A. From this information it may be incorrectly inferred that there exists communication link between C and F. The other links are correctly inferred. In order to eliminate such effects, multiple traceroutes need to be conducted from A to H and the statistically most probable results should be retained.

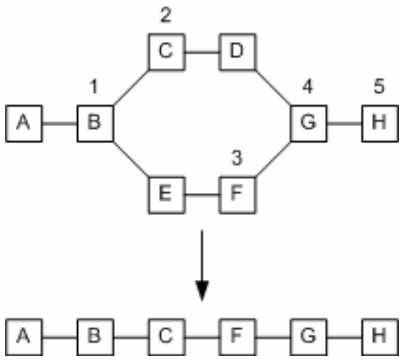

Figure 1: Example of incorrect link inference

Typically, a traceroute would be conducted against each node that is discovered through the host discovery process. The information that is gathered could contain a lot of duplicated data that need to be processed to determine the network topology. This can conveniently be done through a two-phase process. In the first phase, the information is processed to create a set of unique nodes. In the second phase, the information is processed to extract unique links. The extracted information completely describes the network topology.

Figure 2 shows an example that illustrates that some link information might not be determined by conducting the traceroute scans from a single point in a network. In the example, a traceroute is conducted from A to E and from A to F. Following the two-phase process just described, all topological information may be inferred, except the link that exists between C and D. To extract that information a traceroute need to be conducted from either C or E to either D or F. This serves to illustrate that a more complete view of a network is achieved by adding more nodes from which a scan can be conducted, and merging the information from the different scans.

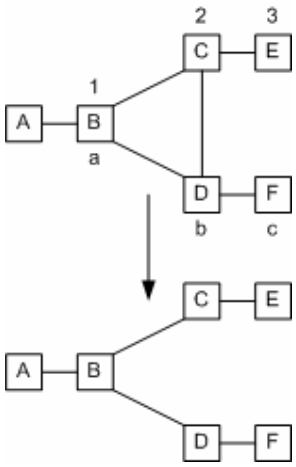

Figure 2: Example of a link that cannot be inferred

5. Usage Determination

Apart from knowing which nodes form part of a network, it is often useful to know how those nodes are being utilised. Techniques for usage determination are discussed in this section.

5.1 Port Scanning

A number of protocols that are used to exchange information between two nodes in a network use the concept of a port. Since multiple programs on a single node could require networked capability at the same time, a method is needed to determine which program to associate with data that arrives at the NIC. The solution is to introduce a port, a unique number that is specified as part of the protocol that contains the data. The port numbers are used by the programs executing on the node.

Table 1 shows a number of well know port numbers used in the Transmission Control Protocol (TCP) and User Datagram Protocol (UDP) range of protocols. A more complete list can be found in RFC1700.

Table 1: Commonly used port numbers

Port Number	Protocol
20, 21	File Transfer Protocol (FTP)
22	Secure Shell (SSH)
23	Telnet Protocol
25	Simple Mail Transfer Protocol (SMTP)
53	Domain Name System (DNS)
67, 68	Dynamic Host Configuration Protocol (DHCP)
80	HyperText Transfer Protocol (HTTP)
110	Post Office Protocol (POP3)
137	NetBIOS Name Service
143	Internet Message Access Protocol (IMAP4)
443	Secure HTTP (HTTPS)
445	Microsoft-DS (Active Directory)

Port scanning is conducted by simply trying to establish a socket connection with a node. The socket connection specifies the IP address of the target node as well as the port number on which communication should take place. If the connection is successful, it means that the port is active on the target machine. If the connection cannot be established, it may be due to the fact that the

port is not open on the target machine, that a communication error occurred, or that the connection is actively refused.

5.2 Service Detection

Port scanning gives an indication of the probable protocols that are being utilised by a node. It is probable in the sense that most commonly used protocols are usually communicating on very specific ports, as indicated in table 1.

However, protocols are being used by programs running on nodes. These programs are called services. A service my implement a protocol to run on a different port than what is expected, or conversely, may run a different protocol than what is expected on a specific port. These exceptions are rare but are a possibility.

In most cases however, a service typically implement a protocol on a specific well defined port. For instance Internet Information Services (IIS) and Apache are both commonly used services that implement the HyperText Transfer Protocol (HTTP) on port 80. It is used to provide web pages to client nodes connecting to it.

Service detection is the process of determining which service (not just the protocol) is running on a specific port. In order to achieve this, the behaviour of the services under certain conditions needs to be known beforehand. For instance, different services may implement the data it embeds in a protocol in a slightly different way, usually in fields of the protocol that does not influence the data detrimentally. These differences define signatures (or a set of behaviours) for the different services.

Service detection starts by port scanning and then testing the open ports against the most likely signatures associated with the port numbers. In the case that the signature cannot be matched, tests are conducted for lesser likely services and even for services that are not expected to be on that port at all.

Version detection takes service detection one step further by aiming to determine the version number of a specific service. For instance, it may be possible to determine whether version 4 or 5 of IIS is communicating on a node by examining the difference between their behaviours.

Nmap is a free tool that can be used for port scanning, service detection, version detection and operating system detection (see next section). It comes with a database containing thousands of signatures of commonly used services.

5.3 Operating System Detection

Operating system detection is accomplished in a similar way as service detection, i.e. through the use of signatures that define the behaviour of the operating system under certain communication conditions.

Through the use of operating system detection, the operating system executing on a node, such as Windows, Linux, Unix or Mac OS can be detected.

6. Visualisation

The techniques of host discovery, topology determination and usage determination are used to create a structural and functional analysis of a computer network. For humans to build insight, it is often necessary to visualise this information. Figure 3 (see previous work by Senekal, 2006) shows an example where such visualisation is used to great effect. As described earlier, the RIRs assign blocks of IP addresses to different organisations. The block of addresses that gets assigned is published, together with the organisation's country. In the example, each network is represented by a dot. These dots are randomly placed (according to a uniform distribution) within

the boundaries of the associated country, creating a computer network density map of the world. From this map it becomes immediately obvious that countries such as the United States, Japan and Europe are more densely populated with computer networks than the rest of the world. Even though the visualisation is not an exact representation of reality, it does facilitate discoveries such as these that would otherwise have been difficult or even impossible to gather from the data. It also enables additional questions to be asked, such as the reason for the apparent lack of computer networks in Africa.

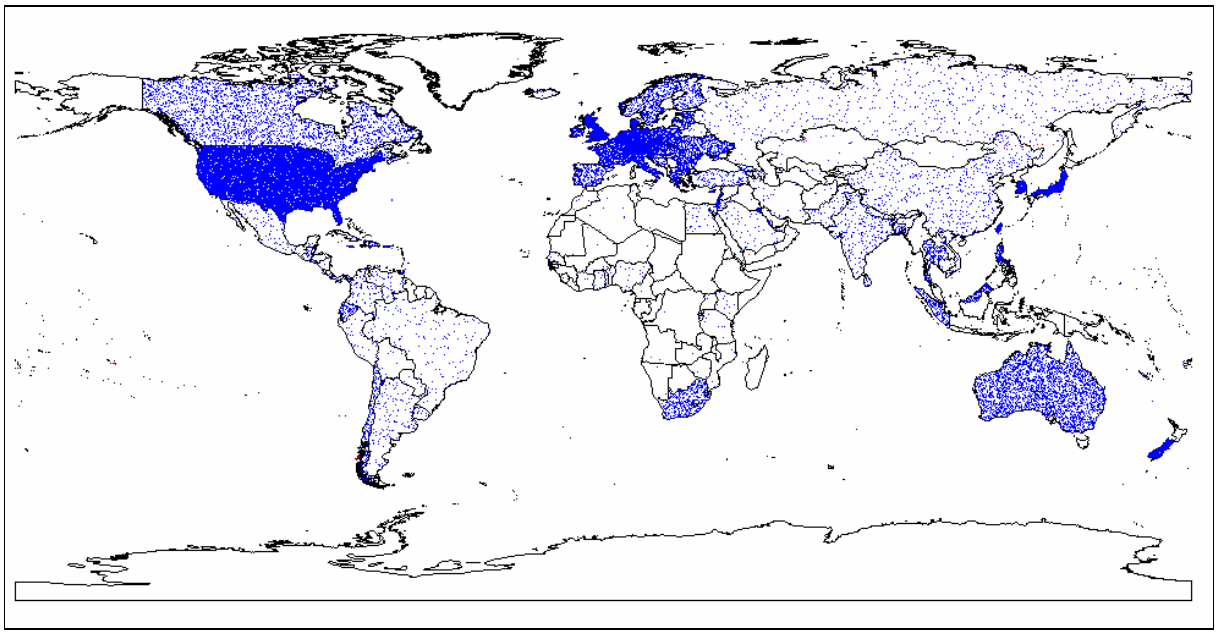

Figure 3: Computer networks density map of the world

6.1 Abstract Layout Algorithms

Abstract layout algorithms can be used to visualise computer networks and other graph topologies. This discussion will concentrate on two-dimensional layouts, since these are commonly used for display on computer monitors, printing, etc. With these techniques, each node is assigned an *(x, y)* coordinate in two-dimensional space with lines between the nodes indicating communication links (or edges, in the more general case of a graph). More complex three-dimensional visualisations are possible and are often found where a user can interact with the visualisation environment.

Figure 4 shows different visualisations of the CSIR's computer network. The data was gathered using the techniques described in earlier sections. About 1500 nodes and a similar amount of communication links were identified. In Figure 4a, a random arrangement of nodes is made within a polygon representing the boundaries of the CSIR's campus. It can clearly be seen that the visualization is not very useful. Figures 4b show an example of a deterministic layout algorithm applied to the data. In a deterministic algorithm, the coordinates are uniquely determined by the topology of the graph. As can be seen, the visualisation creates a layout in which the different internal networks and backbone structures can easily be detected.

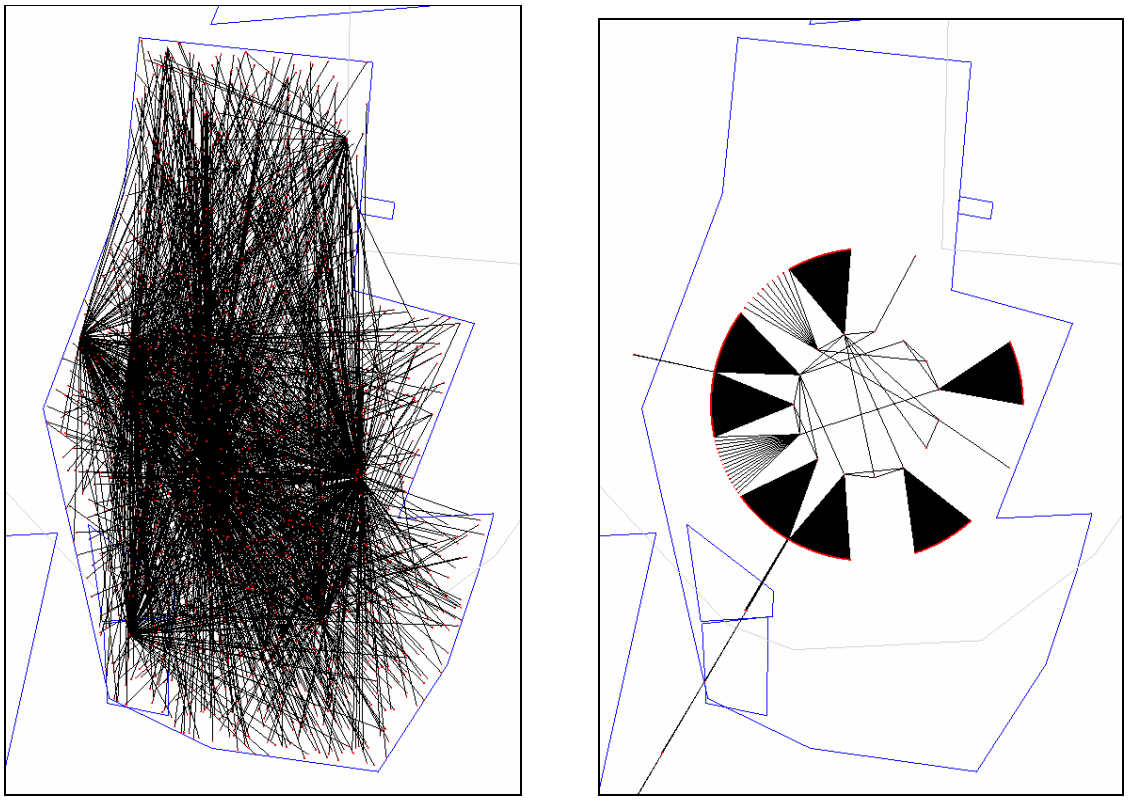

Figure 4: Visualisation using deterministic layout algorithms

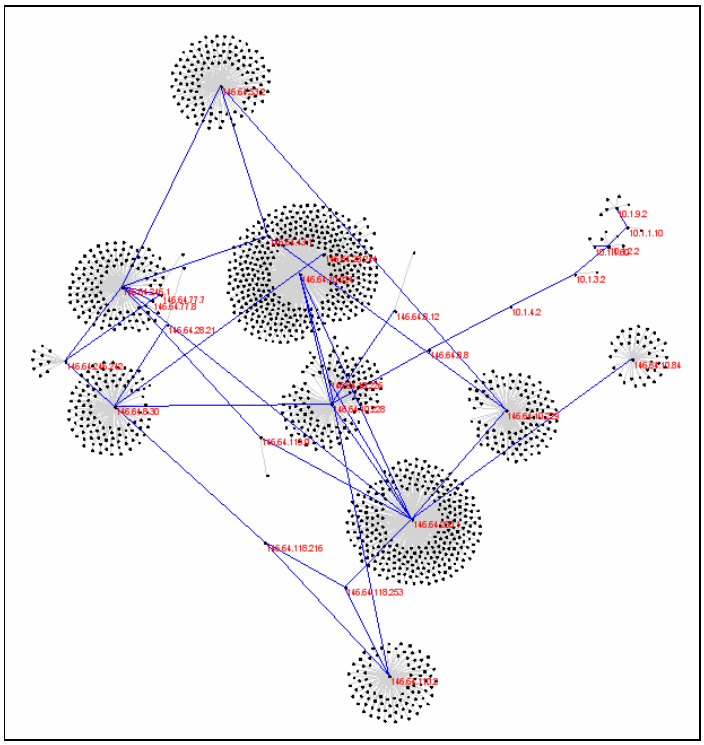

Figure 5: Visualisation using iterative force-directed layout algorithms

Figure 5 shows an example of the output created by an iterative visualisation method. In iterative visualisation methods, changes to the coordinates of each node are made during each iteration until a stable configuration is found. The advantage of iterative methods is that sections of the

network are clustered together, accentuating the underlying structure. In figure 5 a force-directed layout algorithm was used (Churcher, 2004). Each node is modelled as a point charge that repels other nodes in a direction of a line joining the two nodes. The magnitude of the force F with which two nodes repel each other is given by Coulombs law: $|F| = k_C\, q_i\, q_j / r^2$, where k_C is a constant and q_i and q_j are the charge values of the two nodes and r is the distance between the nodes. In physics, k_C is a fixed value. In the force-directed algorithm, k_C is a value that is experimentally determined by the size of the network and q_i and q_j are often assigned values of 1. Nodes that are connected via a link attract each other according to Hooke's law: $|F| = k_H(r - l)$, where k_H is a constant, r is the distance between the nodes and l is the length between the nodes in a rest state. During each iteration, the resultant force on each node is calculated and a change is made to its position based on this force.

6.2 IP Geolocation and Geographic Information Systems

IP geolocation techniques are used to assign geographical coordinates to nodes. This is often used commercially on the Internet, for applications such as target content – where content is displayed on a website based on the area that a user is coming from, fraud detection – where transactions and payments from certain areas or countries are limited, enforcing trade agreements – preventing transactions that would break trade agreements, criminal investigations – to trace back sources of crime, etc.

For military purposes, IP geolocation is a first step if computer network information needs to be visualised using geographical information systems (GIS). An example of GIS software used extensively in the South African National Defence Force (SANDF) is ESRI's ArcGIS software suite. Figure 6 shows an example where a number of backbone nodes in South Africa were geolocated and visualised using Google Inc's free software tool GoogleEarth.

Figure 6: Backbone nodes geolocated and visualised using Google Earth

In database geolocation, a range of IP addresses or even a single IP address is mapped to a set of geographic coordinates. Such a database is created through the analysis of data provided by the

RIRs and other footwork, which is usually a tedious process. Although accurate, the main drawback of database geolocation is the fact that information may become out of date quickly and thus inadequate for military purposes.

An alternative is constraint-based geolocation. A number of nodes from which scans of a network can be conducted, called landmarks or markers, are required. The time it takes for communication to occur between the landmark nodes and a target node is measured. Based on network constraints, such as the expected types of links, the bandwidth available, etc., a sphere (usually projected onto a circle on the earth's surface) around each landmark is created where the target node could possibly be. The intersection of all the spheres indicates the likely position of the target node. The more landmarks are available, the smaller the uncertainty in the position of the target node.

7. Conclusion

Techniques such as host discovery, topology determination, usage determination, network visualisation and IP geolocation can work together to provide a view of the structural and functional components of a large computer network. These techniques are application-oriented and currently in use to map large parts of the Internet.

References

Churcher, N., Irwin, W. and Cook, C. (2004). Inhomogeneous Force-Directed Layout Algorithms in the Visualisation Pipeline: From Layouts to Visualisations, *Proceedings of the 2004 Australasian Symposium on Information Visualisation*, vol. 35, pp. 43-51.

Epp, S.S. (1995). *Discrete Mathematics with Applications*, 2nd ed., Brooks/Cole Publishing Company, Pacific Grove, California.

Huffaker, B., Plummer, D., Moore, D. and Claffy, K.C (2002), Topology discovery by active probing, *Proceedings of the 2002 Symposium on Applications and Internet*.

Hyun Y., Broido A. and Claffy K.C. (2003). Traceroute and BGP as Path Incongruities, *Cooperative Association of Internet Data Analysis – CAIDA*, San Diego Supercomputer Center, University of California, San Diego, USA.

Kurose, J.F. and Ross, K.W. (2001). *Computer Networking – A Top-Down Approach Featuring the Internet*, Addison Wesley Longman Inc.

Lowekamp, B, O'Hallaron, D.R. and Gross, T.R. (2001), Topology Discovery for Large Ethernet Networks, *Proceedings of ACM SIGCOMM 2001*, San Diego, California, pp. 237-248.

Mockapetris, P.V. (1987). Domain Names – Concepts and Facilities, *RFC1034*, http://www.rfc-editor.org/rfc/rfc1034.txt.

Mockapetris, P.V. (1987). Domain Names – Implementation and Specification, *RFC1035*, http://www.rfc-editor.org/rfc/rfc1035.txt.

Plummer, D.C. (1982), An Ethernet Address Resolution Protocol, *RFC826*, http://www.rfc-editor.org/rfc/rfc826.txt.

Postel, J. (1981), Internet Control Message Protocol, *RFC792*, http://www.rfc-editor.org/rfc/rfc792.txt.

Reynolds, J. and Postel (1994), J, Assigned Numbers, RFC1700, http://www.rfc-editor.org/rfc/rfc1700.txt.

Senekal, F.P. (2006), Understanding information warfare through simulation, *CSIR ScienceScope*, vol. 1, no. 2, p. 23.

Siamwalla, R., Sharma, R. and Keshav S. (1999), Discovering Internet Topology, *Proceedings of the IEEE Infocom 1999,* pp. 21-25.

Stallings, W. (2000), *Data & Computer Communications,* 6th ed., Prentice-Hall, New Jersey.