Appendix F
MIMcom VSAT Network

The MIMCOM network was initiated by the Multilateral Initiative on Malaria (MIM), which is an international group coordinating the Malaria research efforts of various organisations.

Since a large proportion of malaria research is done in Africa, a pressing need for these research organisations was proper Internet access, in order to communicate with each other, with funding bodies and with collaborators in the rest of the world. Also, access to journals and the latest publications was needed. It is also much easier publish, when connected to the Internet. At a MIM conference in 1997 in Dakar, Senegal, MIM asked the US National Library of Medicine to lead the project.

The MIMCOM network now includes sites in several African countries, and more are planned. More about the history of the project is available at www.mimcom.net.

1. Introduction

Because the MIMCOM VSAT network is a single network covering sites in several countries, universities might consider it as a model for a joint academic network, or a way to connect several campuses to each other and the Internet.

At every MIMCOM site, access via a local ISP was considered. However, at 11 sites, the only feasible solution was to connect via VSAT. Apart from the 11 VSAT sites, there is connection via a local ISP at a few other locations.

This case study focuses on the MIMCOM VSAT network only. The VSAT network connects sites in Ghana, Kenya, Tanzania, Uganda and Gabon.

The VSAT sites are:

- Kenya: Nairobi (several networks), Kisian, Kilifi and Mbita.
- Ghana: Noguchi and Navrongo
- Tanzania: Amani and Ifakara
- Uganda: Entebbe
- Gabon: Lambarene

At Nairobi, there is also a wireless connection to remote sites. Below is an overview of the network.

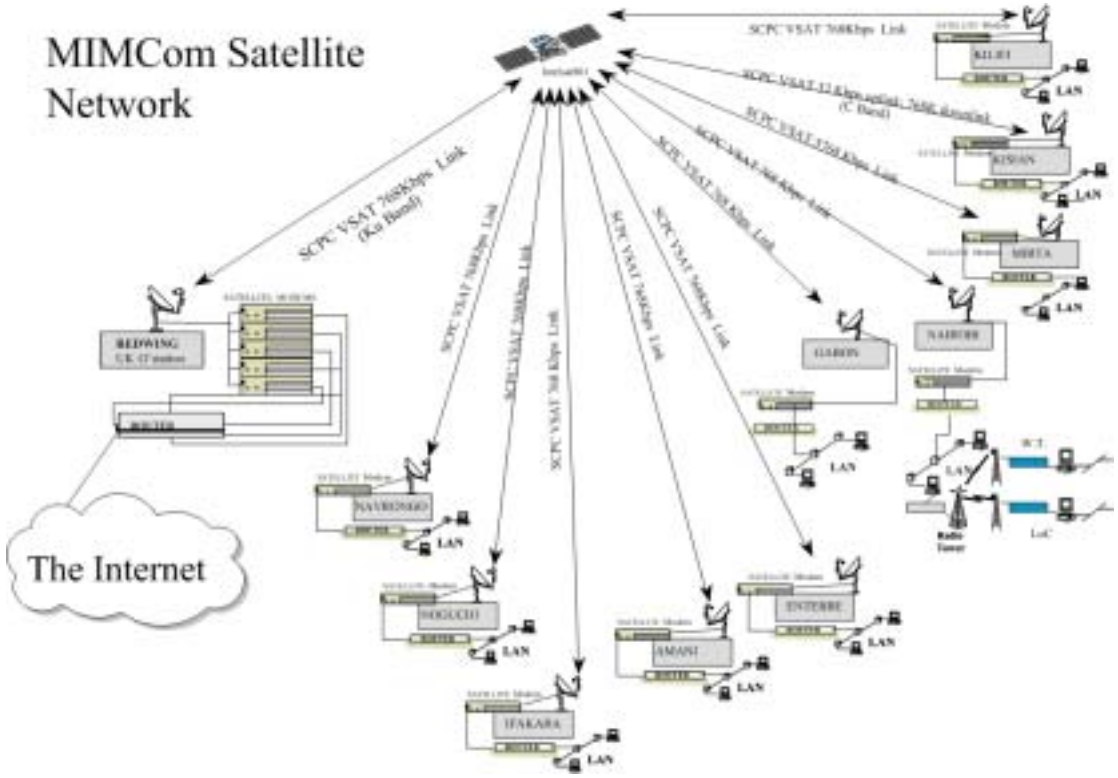

2. The network

The diagram below shows how a site connects to the satellite provider (Redwing Satellite systems near London - see www.redwingsat.com).

Apart from LAN equipment, each site has a Satelcom Megapac router, a satellite modem, and a ChannelMaster 2.4 M dish, as shown in the diagram below.

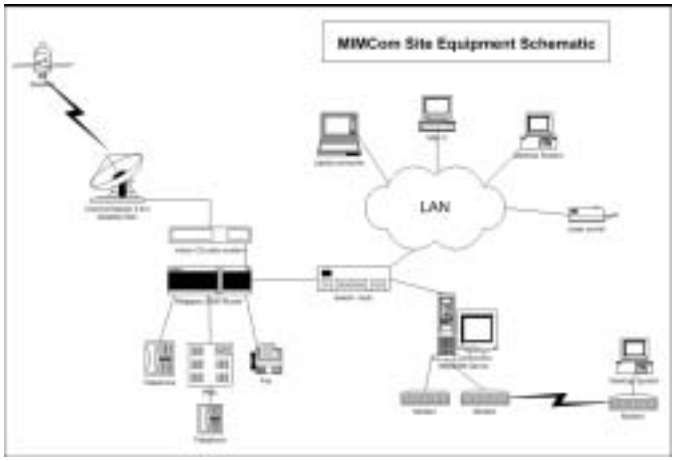

2.1 The telephone system (Voice over Frame Relay)

The Satelcom router has voice ports, which enable communication via Voice over Frame Relay (VoFR). The VoFR system enabled sites to do 3 things:

- Call via satellite with each other for free.
- Free calls to Redwing and Africonnect.
- Enable collaborators in Europe, Japan and the US to call the sites at the price of a call to London.
- If a site requires it, Redwing can activate a breakout to the public telephone system. Then, the site would be able to call other countries at the price of a call from London. However,

since this service would cost some money, or break communications regulations, no site has so far opted for it.

VoFR was found to perform much better than Voice over IP (VoIP). This is logical, because there is less overhead (see below), and frame relay features such as atomisation and prioritisation (both discussed in section 3.5) further improve voice quality.

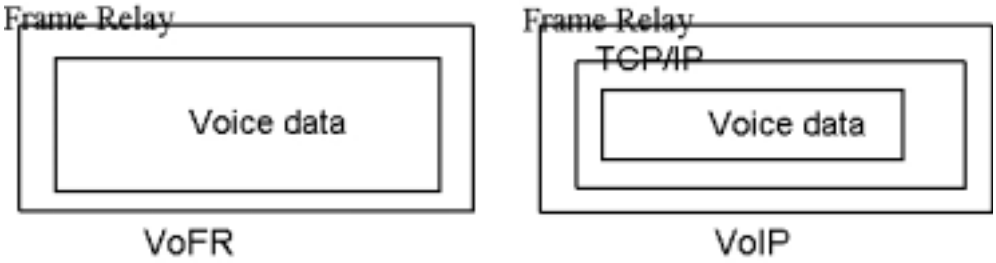

2.2 Data traffic

Frame Relay is the carrier protocol for TCP/IP on the satellite link in the same way as Ethernet is the carrier for TCP/IP on a local area network. Frame Relay is suitable for many different protocols, although the MIMCom network mainly uses TCP/IP and VoFR.

The TCP/IP packets are encapsulated in the frames and each frame contains all the information required to take it across the satellite connection while each packet contains the information required to get it to the final destination.

The frame consists of header information and a data field, followed by a checksum. The header consists of addressing and control information, while the data field usually consists a TCP/IP packet. When they reach the other end, the frames are stripped off and the TCP/IP packets are encapsulated in Ethernet for the next hop, etc as shown below. A TCP/IP packet is similar to a Frame Relay frame, in that it also consists of a header, which contains addressing and control information, a data field and a checksum. All these headers and checksums are considered "overhead", because they are not useful except for getting the data across that hop of the network. The diagram below shows this process of encapsulation, carrying the packets from one type of network to the next.

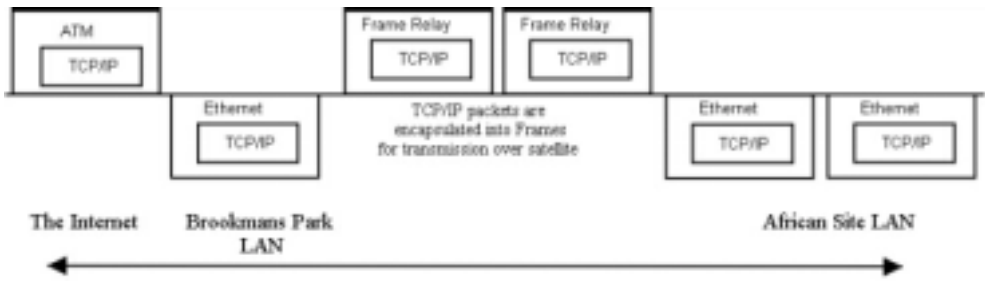

Frames are sent over Virtual Circuits. These circuits are connections between locations on the Frame Relay system. Data Link Connection Identifiers (DLCIs) are used for addressing. A DLCI is a 10-bit address that identifies a particular permanent virtual circuit, a bit like an IP address identifies a device in TCP/IP networks, except that a DLCI identifies a link, while an IP address identifies a device.

The Permanent Virtual Circuits (PVCs) are set up by the carrier that transmits the frame data from A to B. Addressing the data with an address and DLCI will get it over the Frame Relay system to the right address. Communication across a PVC does not require call set-up and termination.

The Switched Virtual Circuits (SVCs) are on-demand connections. They are set up when needed, similar to a telephone call. Both addresses that wish to communicate are assigned a temporary DLCI.

Instead of being allocated a fixed amount of bandwidth for each link, Frame Relay systems allow a Committed Information Rate (CIR) to be guaranteed for each DLCI, and then a Burst Excess (BE) which can be used if other traffic allows it. This means several DLCIs can share a pool of bandwidth while being guaranteed a minimum rate of transmission.

The MegaPAC routers receive all data to be transmitted over the link in TCP/IP packets, and encapsulate the data into frames. These are then addressed to the firewall in London and sent over the satellite. Upon arrival at Brookmans Park they pass through a MegaPAC 2005 router, which converts the frames back into packets.

The diagram below shows the equipment at Brookmans Park in London. The VoFR equipment, and the two Netpilot firewall machines are shown.

The Netpilot (www.netpilot.com), is a Linux based machine that does network address translation, mail forwarding, mail scanning (for viruses) and content filtering. It also acts as a proxy server, and does advertisement blocking. It acts as a firewall, and does port forwarding.

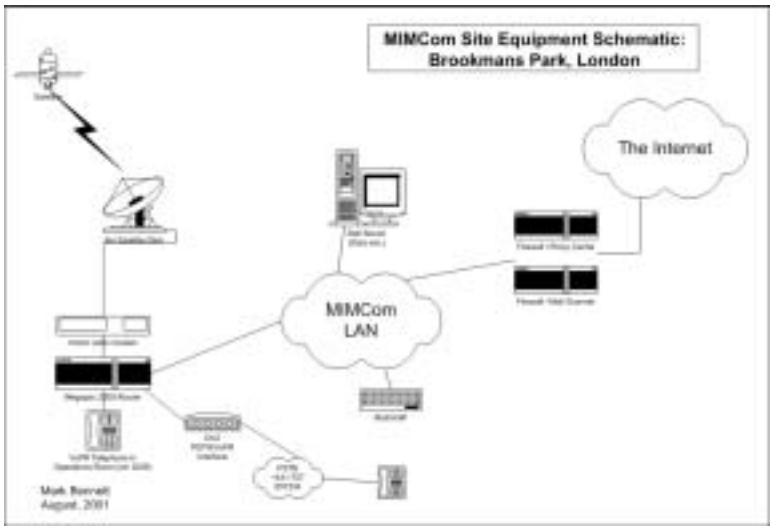

3. Optimisation

Since the MIMCOM network became operational, a number of things have occurred that affected the speed of Internet access – especially web browsing by the end user. These include:

- Considerable expansion of the network, causing pressure on the available bandwidth.
- Growth of the number of users and PCs at the remote sites.
- Introduction of central security and control devices at the satellite ground station in London where the African sites connect to the Internet.

Each of these 3 factors is perceived to have had a negative impact on performance and thus required measures to address them and bring performance back to acceptable levels. What follows are some of the steps taken to identify the problem and provide a remedy.

3.1 The central firewall

In the early stages of the network, the sites had “real world” addresses, and were thus unprotected, except for any local measures they might have put in place. This was not satisfactory, because if something went wrong at a site it could affect other sites. Also, there was no monitoring or content filtering, and nothing was in place to determine what was using the most bandwidth. To this end, a Netpilot firewall was installed. Later, a second Netpilot was installed, because the network grew beyond the performance capabilities of one machine. Since the Netpilot performs caching of Web and FTP traffic, and caching takes place before the traffic gets to the final destination in Africa, it was found this is not the optimal place for caching to be done. The best place for caching that leads to improved response times and slightly lower bandwidth use is at the site. However, content filtering that covers all the sites require proxying at the London ground station. The performance delay that this introduced, is now largely solved by avoiding disk based caching, as described in the next section.

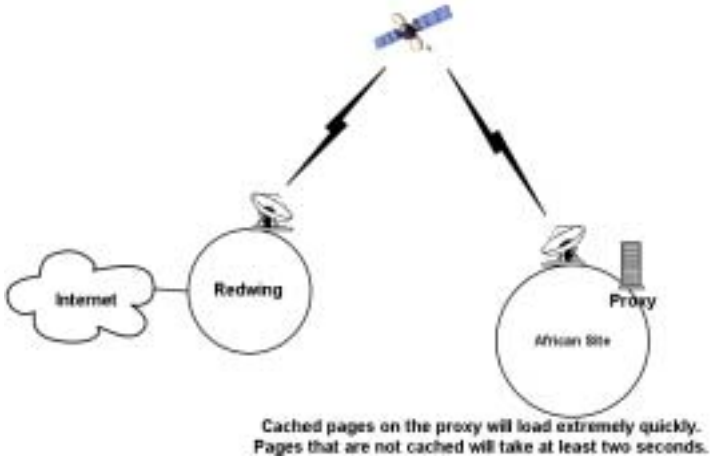

3.1.1 Central Proxy/firewall testing and performance

The Netpilot is a multi-purpose machine that acts as firewall, but also function as a content filtering proxy server, and a virus scanning email relay. It produces log files that can be analysed in order to understand what is going on in the network. Due to complaints after the introduction of the first Netpilot, a number of steps were taken to improve performance. Since the security, content filtering, virus scanning and logs are essential, it could not be removed from the network. A test (see below) was devised to determine the impact of the Netpilot on response times.

The following actions were taken to improve performance. During this time, the tests were running, in order to determine the impact of each action.

- Introduction of a second Netpilot to load balance.
- The memory of both machines were increased from 256 MB to 768MB
- After experimenting with the Squid proxy cache settings, it was found that the best performance is achieved when the amount of disk space allocated to Squid is zero, so that caching takes place only to the memory. This is because disk access is far slower than memory access.
- After these steps did not bring about sufficient improvement, a meeting was arranged with Equiinet (the manufacturer of the Netpilot) to discuss their performance. This step yielded very useful results (see section 3.1.2)

Details of tests

For this test a small website (located in the UK) consisting of typical web pages with graphical content was downloaded every 40 minutes. The total size of the web site was about 2 MB. Downloading a web site was considered a better test of the proxy server than downloading a large file. The performance when downloading a single file was generally good, while the sites complained that viewing a web page was slow.

Three downloads of the same website were compared:

- A download used the Netpilot as proxy. (proxy on)
- A download using the Netpilot as proxy, but not retrieving data from its cache. (proxy on, cache off)
- A download that completely bypassed the Netpilot (proxy off). This was possible because a separate network card was installed for this purpose.

The test was done on a machine at Redwing, and two of the three downloads were forced to go through the Netpilot. The test did not make use of the satellite link because factors such as congestion on the MIMCom bandwidth would have skewed the results. The download times were measured and the results entered into a spreadsheet. The program used to download the web site is called wget. It is part of most versions of Unix / Linux. A Windows version is at http://space.tin.it/computer/hherold. There is also a Windows installer based version at http://homepage.mac.com/shadowboxer/unxutils.exe

Several factors were investigated that could influence proxy performance:

- Content filtering: The Netpilots use N2H2 content filtering, and advertisement – blocking. When a web site is accessed, the Netpilot looks that site up in the N2H2 database to see if the web site is listed under one of the unwanted categories such as pornography, violence, or bandwidth intensive sites such as live video. If it is, the site is blocked. In order to conserve bandwidth, advertisements within a site are also blocked. The concern was whether the content filtering, particularly the database lookup to the N2H2 server, caused an unacceptable delay. To this end, content filtering was turned on some days and off on other days during the test.
- Proxying. To investigate the effect of caching at the satellite ground station, and whether the optimum cache size is used, and whether more memory is needed for caching of web content.

Results

As explained in the section 'Meeting with Equiinet' below, the main factor causing the Netpilots to be bottlenecks were not as expected the N2H2 filtering. The following graph shows the situation as it used to be. The time it took to download a 2MB web site was often more than 5 minutes. The average download time (working hours only) over the period of this graph was
Using proxy: 4 minutes 50 seconds (for both types of downloads via the proxy)
Bypassing proxy: 2 min 9 seconds

Changing parameters such as turning content filtering off did not appear to make any real difference, and poor performance was often observed, as can be seen below. The poor performance seen here was solved after a meeting with Equiinet.

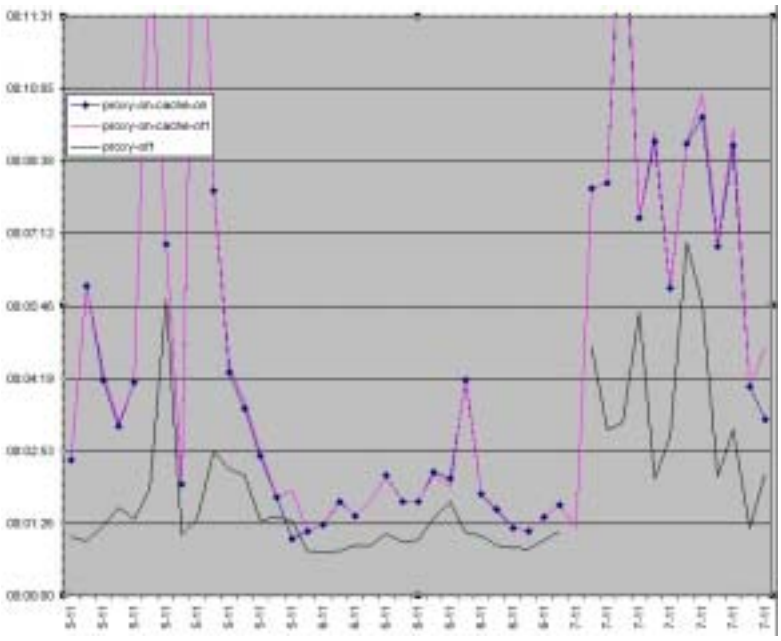

3.1.2 Meeting with Equiinet

Following the meeting with Equiinet, the following points emerged in relation to the tests and changes carried out.

- Email anti-virus scanning is very resource intensive. Equiinet therefore proposed that we split the mail and web functions
- Contrary to expectation, N2H2 filtering might not be the only or even the main problem that causes the Netpilot to slow web access down, because it should take only around 60 ms to validate a URL to the N2H2 server at Bletchley (Milton Keynes)
- The N2H2 list cannot be cached or otherwise kept on the Netpilot, but an N2H2 server can be acquired, and placed on the same LAN as the Netpilot.
- The Netpilot mirrors the contents of one disk to the other once every hour. The purpose of

this mirroring is for backup purposes. When this is taking place it will have a severe impact on performance. However, mirroring can be turned off.

- According to Equiinet, the Netpilot was not designed to be in the centre of a large network. Its specifications are suitable to be the proxy, mail, web server and firewall for a small LAN. (Max 500 users with email & 10% of users browsing at any one time)
- N2H2 is better than competing products & have a database of 20 million sites (other products have 2 million)
- Proxying is a disk bound activity. Changing the cache size to zero appears to be the answer, because then Squid should only cache to its memory, although this leaves the theoretical possibility that Squid will use the page file.
- Proxying can be completely turned off by turning off "Web Access Controls"
- Unfortunately N2H2 requires "Web Access Controls" and proxying.
- Any such custom changes are in danger of being overwritten when the Netpilot software is updated.

The following graph shows the current situation. The differences are that the email and web functions have been separated, the Netpilots had their memory increased from 256MB to 768MB each, disk mirroring was turned off, and the Squid proxy cache size is zero. N2H2 filtering still takes place.

The average download time (working hours only) over the period of this graph was
Using proxy: 1 minutes 50 seconds (for both types of downloads via the proxy)
Bypassing proxy: Security concerns caused this test to be cancelled. It should be the same as before at around 2 min 9 seconds

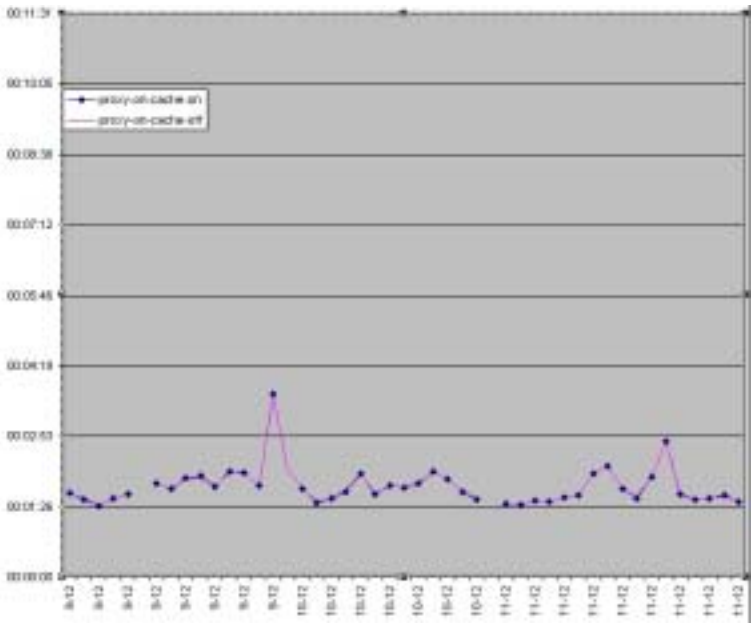

3.1.3 Lessons learned

- Equiinet's suggestions of separating the web proxy function from the email relay and virus scanning function, along with using a Squid cache size of 0, and turning off disk mirroring on the Netpilot that does web proxying appears to have solved the problem, and the web access at the sites is much faster.
- The right place for a proxy server is at the site. The only reason why proxying is required at the satellite ground station is that N2H2 content filtering requires it. Content filtering could be done at the sites, but in practice it will be expensive and difficult to ensure that every site does it, and keeps doing it over time. When cached pages are accessed from a local proxy server, they will reach the user very quickly. Proxying at Brookmans Park in addition to at the sites is an extra step and therefore adds an additional delay for pages that are not cached in the Netpilot's memory. However the above steps make that delay

minimal, because assuming that the proxy service does not use the page file, there is no mechanical device like a hard disk involved (all takes place in memory and CPU), and disk and CPU-intensive applications like virus scanning takes place on a different server.

- The original Netpilot was a low-powered little PC, not aimed at a large network. However, putting enough memory in it, and configuring it to avoid disk-bound operations enabled it to perform well at serving a large network. Nowadays, Equiinet also make Netpilots from more powerful hardware.
- The graph below proves that during the night parts of the web site remains in cache for 40 minutes, which causes the whole website to be downloaded in 31 seconds (blue line). This effect might also occur during the day: commonly accessed web objects might remain in the Netpilot's large memory for a few minutes and will be served very quickly when the next person accesses it.

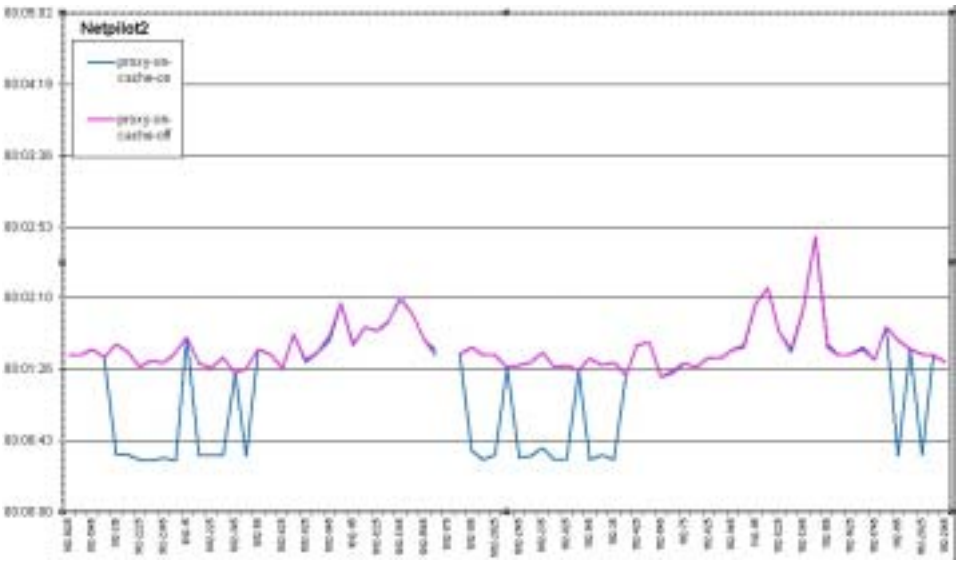

3.2 Proxies at the sites

Initially, no site had a proxy server. The sites were then asked to install a local proxy server. Currently there is a proxy server at every site, but some sections of the site do not make use of it. For example, because CDC Nairobi has its own network and connects via a separate frame relay DLCI, it cannot make use of the Nairobi proxy server. All sites have a local caching DNS, but it is hard to know to what extent the client PCs bypass it, and make use of the Netpilot as their DNS server. (This would make their Internet access very slow).

Site	Proxy	Comment
Kisian	MS ISA server	Kisian recently changed their firewall to force all users to use the local proxy.
Kilifi	Novell Border Manager	This proxy server has recently been revitalised. The system administrators at Kilifi are working to ensure that everyone uses it.
Navrongo	MS Proxy server 2	No problem exists of PCs bypassing the proxy
Noguchi	MS Proxy server 2	No problem exists of PCs bypassing the proxy
Nairobi	MS ISA server	Now works very well – not being bypassed CDCNairobi has no proxy server
Ifakara	MS Proxy server 2	Many PCs are bypassing the proxy server.
Amani	MS Proxy server 2	Many PCs are bypassing the proxy server.
Mbita	MS Proxy server 2	A few PCs are bypassing the proxy server.
Lambarene	MS Proxy server 2	Bypassing of the proxy server not possible
Entebbe	MS Proxy server 2	All the organisations at the site now uses the proxy.

Lessons learned

It is difficult to get a site to install a proxy server after the site rollout. Ideally it should be part of the site installation. When implementing a proxy server one of the major tasks is to

configure the web browser software on each client machine to use the proxy server, and to ensure that this remains so, including on new or repaired PCs. The log files on the Netpilot (the main proxy server at Brookmans park in London) are used to determine at which sites the local proxy servers are being bypassed.

In a set-up such as the one below, the proxy server can easily be bypassed, either by intent or unwittingly - for example if a new computer is introduced to the network and no one remembers to set the proxy settings in Internet Explorer.

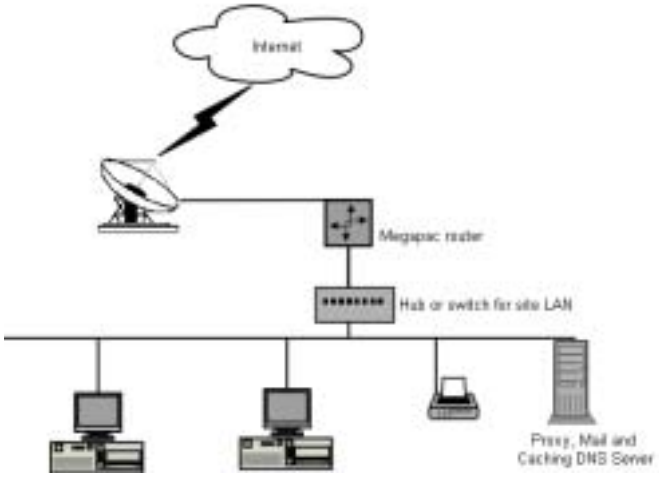

In this case the administrator typically uses one of the following techniques:

Not give out the default gateway address through DCHP. The disadvantage is that some network savvy users who want to bypass the proxy might find or guess the default gateway address.

Using domain or group policies is very useful for configuring the correct proxy server settings for Internet Explorer on all computers in the domain, but is not very useful for preventing the proxy to be bypassed, because it depends on a user logging on to the NT domain. A user with a Windows 95/98/ME computer can cancel his logon, and then bypass the proxy, and someone who knows a local user password of his Windows NT/2000/XP computer, can log on locally and do the same. Domain or group policies are used at some sites, but will not be recommended for the other MIMCom sites.

Since neither of the above methods are reliable, steps were taken to help the system administrators at sites where the proxy is bypassed, to prevent this, using a layout such as the one below, which makes it physically impossible to reach the Internet without going through the proxy server. It works by using 2 network cards; one connected to the LAN and one to the satellite router. This method has been implemented at Noguchi.

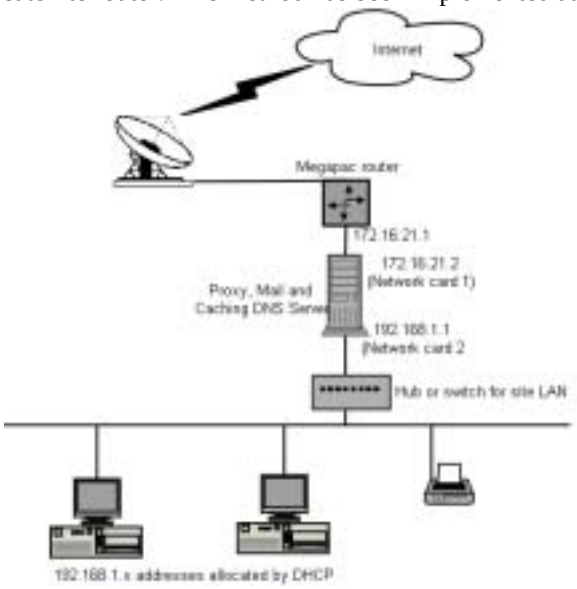

Another reliable way to ensure PCs don't bypass the proxy can be implemented at sites where there is also a local firewall. The firewall can be configured to only allow the proxy server "through", that is to make http requests to the Internet. All other PCs are blocked. In this case, HTTP, HTTPS, FTP, DNS and TCP port 8000 are blocked. (TCP port 8000 can be used to connect to the Netpilot firewall in London as the proxy server for a user's web browser.) This method has been implemented at Kisian.

Some sites have Microsoft ISA Server installed. It has powerful pre-fetching capabilities. A default installation can easily consume more bandwidth than the site has used before, because popular pages with short expiry dates such as news sites, are continually being refreshed. A site with MS ISA Server has to ensure that the bulk of pre-fetching takes place overnight.

3.3 A local caching DNS

Whenever someone accesses a resource on the Internet, a name-to-address lookup is done through the Domain Name Service (DNS). A name lookup from an African site to a DNS server in London across a satellite link takes about one second. This is because of the latency caused by the distance of the satellite to the earth. This represents significant extra time before the first component of the page can even begin loading. Since some elements of the page such as pictures might be retrieved from different places on the Internet, additional lookups occur while the page is loading, each adding more time before the page is fully loaded. This makes web browsing very slow from a human perspective, even if the bandwidth is enough.

When someone visits a website and someone else visits the same site an hour or a week later, the address will typically remain the same. It therefore makes sense to keep a local copy of the TCP/IP addresses of commonly used sites such as www.google.com that everyone can use. When using a local caching DNS, the time name lookups usually take will be negligible in terms of user perception.

Although a single name lookup does not use much bandwidth, so many lookups are required for each web user that significant savings can be made by implementing a local caching DNS server. The main benefit would be faster web access.

Effect on response times: Using a local caching DNS server will improve the average load time of a page by 1 second or more, depending on the complexity of the page.

Effect on bandwidth usage: The effect will be a very small reduction in bandwidth usage from which all sites could benefit. There is no way at the moment to measure the size of this reduction.

3.4 Other factors influencing performance

Rain, sunspots and other customers interfering on the satellite frequency are factors that occasionally cause temporary problems. Nothing can be done about them, except releasing the sunspot chart that Redwing makes available to the sites. User factors such as large downloads, email forwarding loops and web based email, are possible to deal with.

3.4.1 Large downloads and email forwarding loops

Occasionally, a single user making a mistake can cause a problem. For example a user whose MIMCOM account was configured to forward all mail to his Yahoo account went on holiday. All emails sent to him in his absence were still forwarded to his Yahoo account, which can only grow to 2 MB. When the Yahoo account got full, it started bouncing the emails back to the MIMCOM account, which immediately forwarded the bounce message back to the Yahoo account. An email loop was formed that sent hundreds of thousands of emails back and forth, generating massive traffic and crashing the Netpilot.

A user may also start several simultaneous downloads, or download enormous files such as 650MB ISO images. The only solution to this kind of problem is that the system administrator at every site should keep an eye on the site's bandwidth usage, and they train and inform their users as best as possible.

3.4.2 Web based email

The figure below (created by Sawmill) indicates the amount of bandwidth spent on web-based email on 6 January 2003. The total bandwidth use for web browsing was 2126.68 MB. Of this,

363.22 MB was used for web-based email. Therefore, 17 % of all web based traffic was to web based email sites

	Host or domain	Page views	Bandwidth	Visitors	Page views bar
1	172.16.11.200	3,716	57.13M	1	
2	172.16.15.2	3,381	53.16M	1	
3	172.16.13.3	2,202	41.31M	1	
4	172.16.18.3	2,350	30.64M	1	
5	172.16.16.3	1,605	21.15M	1	
6	172.16.17.3	798	10.35M	1	
7	172.16.15.95	271	7.94M	1	
8	172.16.12.3	353	6.92M	1	
9	172.16.14.3	541	6.88M	1	
10	172.16.19.4	559	6.80M	1	
	141 other items	7,047	120.94M	-	
	Total	23,823	363.22M		

more rows | fewer rows | 10 rows | all rows | export table | export all

Web based email has an adverse effect on performance because not only does web based email use vastly more bandwidth than regular email:

- Web based email pages contain graphically intensive advertisements.
- Web based email services such as Hotmail are prime targets of spam - on average, 80% of email in Hotmail is spam. Dealing with that requires bandwidth.
- Bandwidth is used every time a user looks at the same email, whereas regular email uses the bandwidth only once to get an email.
- Even local email messages travel via the satellite link to the USA and back across the satellite link to a colleague in the same location.

The second problem with web-based email is that in the MIMCom network, there is a prioritisation system that delays email to allow web traffic through. This is one of the optimisations that help to make the best use of the available bandwidth. Delaying emails by a couple of seconds when the network is busy is not a problem when that speeds up the load time of web pages. Users of web-based email undermine this arrangement, because they get access to their email at web priority. This is bad for all other users of the system, especially when attachments are sent.

Some sites such as Navrongo were using Hotmail almost exclusively, while at other sites there was a mixture of web mail and traditional email usage. At Noguchi web mail is blocked. Sites were asked to limit web-based email.

3.4.3 Keeping in control

- In order to cut down unnecessary usage, it is important to keep informed about how the system is used. To this end, the various statistics systems have been made available, and the sysops (system administrators) at all the sites were given access to them. Furthermore, training was given, and usage statistics for each site is emailed to the relevant sysop every week, to enable him to keep informed about bandwidth usage at his site, and to cut down on inappropriate usage.
- In order to further reduce bandwidth that is used for inappropriate purposes, acceptable use policies have been implemented at some sites, and others are being encouraged to do the same.
- Since Windows updates consume such a large amount of bandwidth, sites are encouraged to either disable Windows updates, or utilise Microsoft's Software Update Server (which is free). The value of these updates for workstations in the MIMCOM network is debatable, because the updates are mainly for security, and the majority of users in the MIMCOM network are protected by two firewalls. Microsoft Software Update Server can be installed on the same machine as the proxy server. Small sites are encouraged to disable updates.

Sites were encouraged to add content filtering to their proxy servers - using products such as Websense - to increase the amount of inappropriate material that is blocked and refine their local control and statistics

3.5 Router configuration

The Satelcom Megapac routers are high performance machines with several features that enable the network to make better use of the available bandwidth.

3.5.1 Atomisation

When there is a VoFR conversation in progress, the Megapac routers make the frame sizes smaller to ensure that frames containing voice traffic do not get stuck behind large frames that are in the process of being transmitted. This process is called atomisation. It works very well for getting better quality voice traffic. However, it has some implications for TCP/IP traffic. The TCP/IP packet sizes have to become smaller, and when the call completed and frame size returns to normal, TCP/IP packet sizes will adapt as well.

3.5.2 Prioritisation and mail separation.

A Megapac router does not read the TCP/IP packet to determine which kind of TCP/IP protocol it is (e.g. mail, ftp or web), but it does read the IP address destination from the TCP/IP header of the packet within the data section of the frame. This enables it to separate different kinds of traffic based on the destination TCP/IP address. Additionally, it is able to determine if it is non-TCP/IP traffic such as VoFR. Based on this information, it can prioritise in the order of sensitivity to latency. For the inbound traffic (from the site to London) it works as follows:

- VoFR traffic gets the highest priority because if the frames don't arrive in time no conversation is possible. The prioritisation is achieved by sending voice over a separate DLCI and giving that DLCI a higher priority. Prioritisation means when there are different kinds of traffic in the router buffer, the higher priority traffic will be sent first.
- All other traffic except email is on a lower priority DLCI.
- Email traffic is also on a separate DLCI. It is currently not given a lower priority, but is limited by giving it a CIR which is a proportion of the site's overall outbound CIR (from London), and which can burst to the full bandwidth available to the site. This means that for the inbound link, email traffic will not delay web requests because it can only use more than a certain percentage of the bandwidth if it is not in use by voice or web traffic. This effectively deprioritise mail traffic because it seldom matters if an email arrives a few seconds later. The Megapac routers identify mail traffic as everything addressed to the TCP/IP address of the mail relay server.

For the outbound traffic (from London to the site) it works as follows:

- Outbound VoFR traffic is also prioritised by sending voice over a separate DLCI and giving that DLCI a higher priority.
- All other traffic except email is on a lower priority DLCI
- Separation of email traffic will soon be implemented; using the TCP/IP addresses of the mail servers at the African sites as identifiers. For sites where the mail server is behind a local firewall, the address of that firewall can be specified when setting up the DLCI.

3.5.3 Future developments

Future improvements that can be achieved using the Satelcom Megapac routers include multicast, compression and encryption. Multicast means that since all traffic from London to the sites are broadcast to all sites (each site just accepting the traffic addressed to it), commonly accessed pages and other web elements can be cached at the other sites when they have been sent to the first site that requested them. When someone requests them from another site it is immediately and locally available. It will be a long time before this can be implemented, because it will require additional hardware at each site.

Compression and encryption can soon be implemented. Compression could lead to a slightly better use of the available bandwidth. Encryption will make transmission more secure. The VSAT protocol, which is a new feature of the MegaPAC software, contains features such as TCP acknowledgement spoofing to improve throughput.

4. Deciding about bandwidth

The table below shows the bandwidth allocation before it was increased. The total inbound bandwidth (from the African sites into London) was 608 Kbps and the outbound from London to the African sites was 768 Kbps. (The terms inbound and outbound are from the London (ISP) perspective throughout this document).

Site	CIR In	CIR Out	BE	Current maximum possible outbound (CIR Out+ BE)
Kisian	160	128	512	640
Kilifi	64	64	256	320
Navrongo	32	32	256	288
Noguchi	32	32	256	288
Nairobi	64	64	512	576
Ifakara	32	32	192	224
Amani	32	32	192	224
Mbita	32	32	192	224
Entebbe	64	64	512	576
CDC Nairobi	64	64	512	576
Gabon	32	32	256	288
TOTAL	608	576		

Each site has an inbound committed information rate (CIR) allocation, which means that there is no contention with other sites on the inbound link. Each site also has an outbound committed information rate (CIR) and an outbound burst excess (BE). The BE is a pool of additional bandwidth that can be used by any site on a contended basis. It consists mostly of unused bandwidth from other sites. This extra pool of bandwidth has the advantage that whenever the other sites are not using it, any site can get access to high bandwidth, enabling them to complete a large download quickly, for example. Typical Internet usage is very spiky, because someone might do a large download, and afterwards use that file and consume no bandwidth for the rest of the day. The shared bandwidth model is the most economical way to deal with spiky traffic, not least because the sites are in different time zones, so some sites will get good access to the BE because at other sites where the working day has ended, or is yet to begin, the bandwidth usage will be low. Also, large object requests will usually be spread out randomly. Having access to more bandwidth causes each large object download to finish far quicker than it would have done if the site had only access to the amount of CIR it can afford.

The burst excess consists of 192 Kbps that was always available and the rest is taken from the sites' outbound CIR. However, this does not mean that there is ever contention on the CIR bandwidth due to the Burst Excess using unused bandwidth from other sites. The frame relay router allocates the CIR first in each clock cycle. If not all the CIR is used, the rest is available for BE for any site to use, up to the maximum of that site's allocation. Because bandwidth is allocated anew in every clock cycle, there is not even a delay before a site can to "reclaim" its bandwidth.

Therefore, this model of shared bandwidth is not the same as a contended system, because every site has full access to its CIR, and high usage at other sites can never rob a site of its CIR.

Allowing all sites to burst to the full 768 Kbps and doing away with the 192 Kbps that is dedicated to BE will improve the service during periods of low usage. It is possible that making this change will improve or at least be neutral to the performance at all times, and benefit performance during times of low usage.

4.1 User profile and human factors.

The users in the MIMCOM network are a mixture of researchers, admin and general staff, system administrators, and users of libraries or open areas. This last category includes for

example students using the library at KEMRI (Kenya Medical Research Institute). Staff members are generally allowed casual use after hours. This is a measure to retain staff at remote locations, where there is little to do after hours, and to enable them to keep in touch with their families and news sites from their countries. By far the biggest users are the system administrators, who download large files such as anti-virus updates and service packs every day.

The following information has been gathered for Noguchi over a period of 14 days

Total number of authenticated users	122
Total unique hosts	71
Average sessions per day	127
Average page views per session	46
Median visits per visitor: 23.00	23
Average visits per visitor: 25.11	25
Average session duration:	00:23:15
Total Noguchi web traffic over this period	2613.85 MB
Total MIMCOM web traffic over this period	22630 MB

If these figures are extrapolated to calculate the possible total number of users in the MIMCOM network, it works out at 1056. (This would only be a good indication of the actual number if the user and usage profiles at the other sites were similar to Noguchi. The exact number of users in the MIMCOM network is not known, because the network keeps growing, and it is hard for an external company to keep figures up to date.)

Human factors such as acceptable load time have been studied many times. The results of three studies by Anna Bouch (University College - London), Allan Kuchinsky and Nina Bhatti (Hewlett Packard Labs - Palo Alto) on how long users are prepared to wait will be discussed here. A more detailed discussion can be viewed at
http://www.humanfactors.com/downloads/apr012.htm.

In their first experiment, users rated the load time for a web page as "good" when it took up to 5 seconds, "average" when it took from 6 to 10 seconds and "poor" when it took more than 10 seconds.

In their second experiment, users could press a button labelled "Increase Quality" when they felt the load time was slow. The average time before pressing the "Increase Quality" button was 8.6 seconds.

In the third study, they found that web pages that loaded incrementally with the banner first, text next and graphics last. "Under these conditions, users were much more tolerant of longer latencies. The test subjects rated the delay as "good" with latencies up to 39 seconds, and "poor" for those over 56 seconds."

Since proxying defeats the incremental loading of pages, the MIMCom network should probably attempt to deliver pages that load below 10 seconds. In practice some pages that are delivered from the local proxy server, or from the memory cache on the Netpilot will often load much faster, whereas at other times congestion on both the inbound and outbound links, and as well as the latency inherent in a satellite connection, will combine to cause a poor load time.

The question is how do these factors affect the MIMCom users. Due to limited contact with users, the only clue may be in session data from that is calculated from the Proxy server logs, using Sawmill. The average session time for a user on the Internet is about 32 minutes (http://www.nielsen-netratings.com/hot_off_the_net_i.jsp), which is slightly longer than what is seen on the MIMCOM network (00:23:15). The fact that it is lower may be a result of relatively strict usage policy at Noguchi; less non-work access is allowed compared to the other sites, and possibly less than the norm on the Internet. The table below shows how average web session length could have been affected by the bandwidth changes. It suggests that higher bandwidth might cause users to use the Internet slightly longer, reflecting the fact they are frustrated less.

Other factors that might have caused the differences were not taken into consideration. In general it can be expected that a better service will cause users to use the Internet more up to a point, and therefore cause higher usage.

Bandwidth level	Average session duration	Comment
CIR-in-32 CIR-Out-32 BE-256	00:24:47	Original bandwidth level.
CIR-in-64 CIR-Out-32 BE-256	00:23:31	
CIR-in-48 CIR-Out-32 BE-256	00:24:22	
CIR-in-64 CIR-out-192 BE-0	00:22:55	
CIR-in-64 CIR-out-128 BE-256	00:27:16	

4.2 Statistics

The statistics system indicates that the broadcast utilisation (TX) was near 100% for most of the working day. The graph below shows very clearly that there is not enough total bandwidth for the amount of usage.

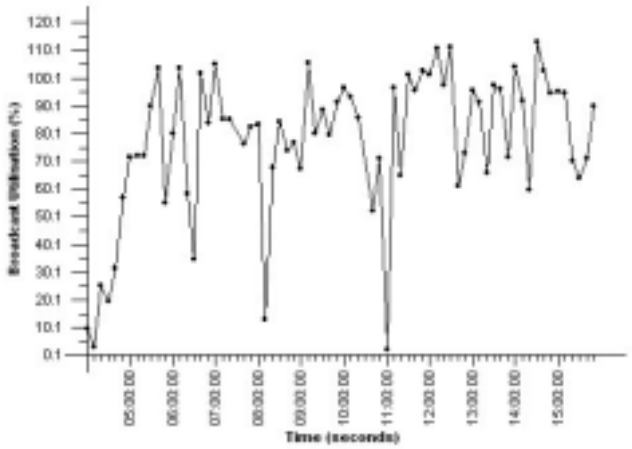

Another example (below) shows Navrongo's TX utilisation and overall TX. The overall TX utilisation is near 100%. There are areas of the graph where high Navrongo TX utilisation corresponds with lower overall TX. When the red line is at 100% flat topping occurs. That is why the majority of Navrongo's peaks are when the TX is not 100%. Navrongo peaks to 300% whenever the TX utilisation allows for it.

Percentages of more than 100% is possible because when it reaches 100% of its available bandwidth and needs more, a site makes use of Burst Excess (BE), which is transmit bandwidth (outbound from London) not being used by other sites at that point in time. Theoretically, Navrongo can reach more than 300%, but there is just never enough BE left by the other sites for this to happen. The maximum utilisation Navrongo can achieve is (CIR+BE)/CIR)*100. This is (32+256)/32 or 900%, so Navrongo is not flat topping, or does not need that level, but the total bandwidth is "flat topping".

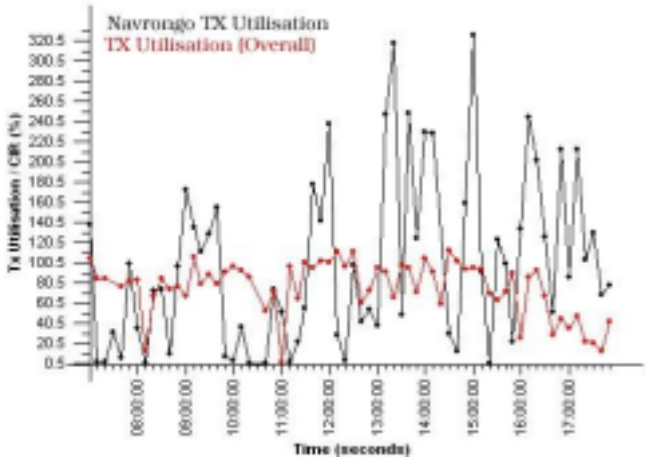

The high overall TX utilisation is a limiting factor for Navrongo to get a share of the Burst Excess.

4.3 How much bandwidth is needed?

In a network of over 1000 users and a total bandwidth of only 768 Kbps, congestion on is predictable, because there is only 0.768 Kbps per user. For comparison, when connected via modem dialup connection, a user has 56 Kbps available to the ISP's point of presence, although the bandwidth from there might be contended. Nonetheless a bit rate of 4KB (kilobytes) or more is achievable, which equates to more than 32 Kbps. (kilobits per second). Typical Internet use such as web browsing and email is relatively slow, but still acceptable. A comparison with a modem link is of course entirely misleading, because not all the users in the MIMCOM network use the Internet at the same time. Additionally, not all connected users use bandwidth all the time, because in between requesting web pages they are also reading them. Another factor is that the higher the available bandwidth the less is needed per user because when someone does use bandwidth, his request can be satisfied faster, leaving him with something to read and the "big pipe to the Internet" open for the next user. The final reason why the comparison with a modem link would be wrong is that most of the sites have local proxy servers and local caching DNS, which mean that a significant percentage of their web requests can be served from a local server. However, a faster link can also cause users to use the Internet more intensively than they would otherwise have done.

The graph of Bandwidth required (and bandwidth required per user) against number of users can be expected to look something like this:

Bandwidth required and bandwidth required per user against number of users.

4.3.1 Contention ratio's

A better comparison than with a single user dialup link would be to look at other bandwidth per user figures. For example, the Ventura Unified School District in the USA had 2 Kbps per user in 2001 and planned to upgrade to 20 Kbps per user. This document can be viewed at http://www.ventura.k12.ca.us/vusd/HTMLobj-1270/a_VUSD_Tech_Plan_12-11-2001.doc. Other such comparisons are difficult to obtain.

Clearly, 20 Kbps per user is not affordable in a satellite network, and the school authority is perhaps only upgrading to that level because bandwidth in the USA is very cheap. From Proxy server sizing guidelines we see that 7.5 Kbps per user for about 100 users is considered sufficient for web, FTP and streaming multimedia. For a larger user base of 1000 users where most users don't use bandwidth intensive applications like streaming media this figure would be correspondingly lower, perhaps even as low as between 1 and 2 Kbps per user. This figure is made more realistic by the use of local proxy servers and local caching DNS servers as described in the Site Proxy section below.

Hughes Network Systems make the following assumptions for their DirecPC satellite Internet service:

- 175 subscribers per 128 Kbps. This is the same as the MIMCOM network, which currently has between 1000 and 1200 users. This gives 0.73 Kbps per user. The Hughes network can use such a low bandwidth required per user because it has so many users (See figure 6 above)
- 10% of subscribers will be logged on, and 5% of logged on users will be active – actually using bandwidth. This is not a good assumption for the MIMCOM network where more than 110 users are active during peak times, and some sites additionally make large data

transfers.

4.3.2 Proposed optimum bandwidth for each site

The proposed bandwidth changes, made necessary by the rapid increase of users, do away with the limitations on burst excess, allowing each site to burst to the full amount of available bandwidth. This should be better because it will be possible to satisfy a large object request quicker, leaving the connection open sooner. The 192 Kbps dedicated to BE is also dropped, because there is no point in refusing a site access to some bandwidth when no other site is using it. These figures are still too low to fully solve the problem and take future growth into account, but is also based on what the sites can afford. Since this proposal, only a few sites have increased their bandwidth.

Site	CIR In		CIR Out		BE		Total Out	
	Current	Proposed	Current	Proposed	Current	Proposed	Current	Proposed
Kisian	160	160	128	224	512	864	640	1088
Kilifi	64	64	64	160	256	928	320	1088
Navrongo	32	64	32	64	256	1024	288	1088
Noguchi	32	64	32	96	256	992	288	1088
Nairobi	64	64	64	192	512	896	576	1088
Ifakara	32	32	32	64	192	1024	224	1088
Amani	32	32	32	32	192	1056	224	1088
Mbita	32	32	32	64	192	1024	224	1088
Entebbe	64	96	64	96	512	992	576	1088
CDC Nairobi	64	64	64	64	512	1024	576	1088
Gabon	32	32	32	32	256	1056	288	1088
TOTAL	608	704	576	1088				

5. Conclusion

The MIMCOM network is a possible model for connectivity for a university with several remote campuses, or even for several educational institutions

The MIMCOM network is really a cooperative organisation, consisting of fully independent organisations. Each research organisation contributes towards its own bandwidth costs. There is a central function for someone to make decisions such as how to improve the central firewalls, decide how the sites should pay, how much bandwidth is needed and the way BE is allocated. This central function is currently shared between Redwing, NLM and Africonnect in an ill-defined way. Since the particpating organisations are entirely independent, no central decision such as requiring the sites to ban web based email can be enforced. Despite this, the sites do tend to do what is needed (such as installing a local proxy server.)

Future development of the network might require sites to sign up to a usage agreement, and appointing an entity to do the central function. This entity might be an office based in Africa, whose task it will also be to evaluate the service and prices obtained from the ISP (Redwing)