Bypassing the Home Computing Bottleneck: The Suburban Area Network

Carlo Kopp and Ronald Pose

Department of Computer Science,
Monash University,
Clayton, Victoria 3168,
AUSTRALIA
E-mail: {carlo,rdp}@cs.monash.edu.au

Abstract. The Internet has become an essential tool in many professions. Also now common is the home office and people working from home. What is missing from the home environment is the high-speed computer networking that one finds at work. Commercial services which enable high bandwidth Internet connections are quite expensive and cannot usually be justified. This paper presents a new cooperative computer network architecture which promises economical, high bandwidth computer internetworking for a community of cooperating users, such as may be found in a university or a company branch office. The Suburban Area Network involves both the physical links and a cooperative routing scheme which provides the required quality of service. The aim is for very low cost implementation. A pilot study of the techniques outlined herein is to be undertaken at Monash University.

1 Introduction

The rapid growth of the Internet and the ever increasing demand of bandwidth to home users of such services has exposed a serious bottleneck in the existing data distribution architecture.

The Internet has become an essential tool in many professions. Also now common is the home office and people working from home. What is missing from the home environment is the high-speed computer networking that one finds at work. Commercial services which enable high bandwidth Internet connections are quite expensive and cannot usually be justified. This paper presents a new cooperative computer network architecture which promises economical, high bandwidth computer internetworking for a community of cooperating users, such as may be found in a university or a company branch office. The Suburban Area Network involves both the physical links and a cooperative routing scheme which provides the required quality of service. The aim is for very low cost implementation. A pilot study of the techniques outlined herein is to be undertaken at Monash University.

Ideally a home user should be able to enjoy similar bandwidth in the household, as the user receives at work. Most workplaces today provide a full 10 or 100 Mbit/s bandwidth to on-site server hosts, and typically between 64 kbit/s and several Mbit/s access speed to off-site server hosts.

It may be argued that in Australia higher bandwidth between local sites may be of limited value for activities such as web browsing. This, it may be argued, is because Australia's primary overseas links constitute such a performance bottleneck, as to make off-site bandwidth beyond 2 Mbit/s of limited utility.

Traffic caching proxy web servers at a site's topological entry point to the outside network can alleviate some of the performance losses seen when accessing server hosts other than those local to the site.

The home user today has very limited options in terms of cost effective medium to high speed connectivity. Most domestic users will employ a data-compressing analogue voiceband serial modem, providing effective bandwidth for PPP (Point-Point-Protocol as per RFC1661, 1548) [RFC1661], or SLIP (Serial Line IP as per RFC1055) [RFC1055] traffic somewhere between 9600 bit/s and 56 kbit/s. Some home users will employ a 64 kbit/s ISDN-B channel (Integrated Services Digital Network), but given the higher cost of ISDN compared to analogue POTS (Plain Old Telephone Service) services, this will most likely only be the case with home users who are subsidised by an employer.

A bandwidth of tens of kilobits per second is of limited utility today. It is clearly adequate for ASCII character mode services such as remote login to hosts, email and smaller file transfers, and browsing web sites with little graphics. It is not adequate by any means for more sophisticated services such as NFS (Network File System) {RFC1813] or CIFS (Common Internet File System, formerly Server Message Block or SMB protocol) file serving, running X protocol to home sites, transferring large files by FTP and browsing graphics intensive web sites. With the high demand for the latter service today, it is evident that the voiceband POTS modem and the ISDN-B are inadequate and for all practical purposes have been made obsolete.

Two new services should become available to household users over the next decade. Both offer important gains in performance compared to POTS and ISDN-B, but both also have limitations and are likely to be much too costly for wider use.

2 Cable Modems

Cable modem technology [Kopp, 1996a; Kopp, 1996b] is designed to exploit the considerable bandwidth available in the cable television distribution medium. Cable TV utilises a 75 Ohm wideband coaxial cable for distribution of services to end user premises. Subject to cable bandwidth and coupling transformer bandwidth performance over one hundred 6 MHz analogue slots may be available to carry either TV broadcast channels, or data. A tree-like distribution topology is employed, with linear cable segments tapped with spurs to individual user premises. Figure 1 depicts the structure and deployment of cable modems and their spectrum allocation.

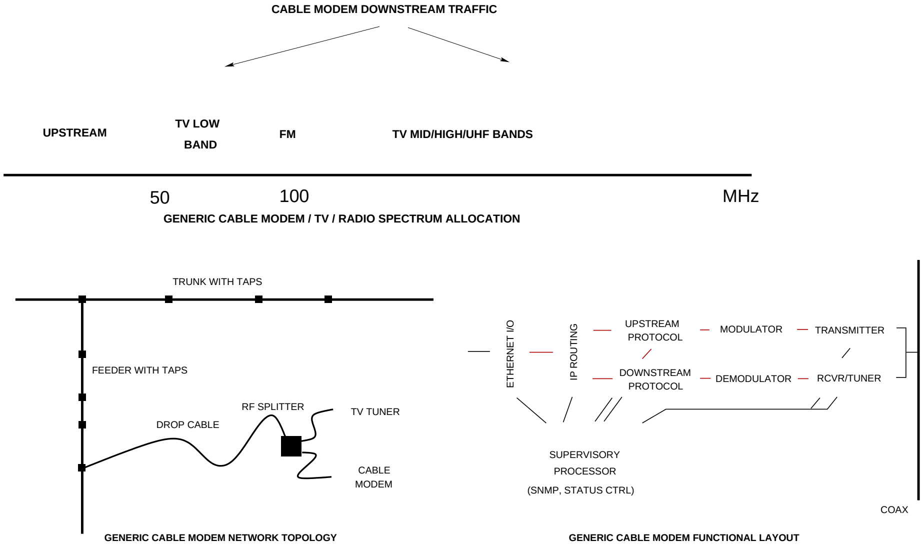

Figure 1. Cable Modem

The paradigm adopted by most cable modem vendors at this time is largely borrowed from the Ethernet/Aloha model. Where it differs fundamentally is that it is asymmetric, and separate bands or channels are used for the user-to-network (upstream) path and the network-to-user (downstream) path.

The downstream channel will typically utilise a single 6 MHz TV channel, shared between tens to hundreds of users. Modulation schemes vary between vendors, QPSK (Quadrature Phase Shift Keying) and particularly QAM (Quadrature Amplitude Modulation) schemes are the most popular. Available downstream channel bandwidths are quoted between hundreds of kilobits/s to 36 Megabits/s. Data packets using formats analogous to Ethernet are thus broadcast to all stations downstream, in a manner analogous to Ethernet transmission.

Upstream communications in cable modem schemes usually operate at carrier frequencies below 50 MHz, most commonly using QPSK modulation schemes for increased interference rejection. Upstream bandwidths vary between tens of kilobits/s up to several Megabits/sec.

At this time the IEEE 802.14 standard is still in definition, it is expected to provide 27 Megabits/s downstream using QAM-64 and between 1.5 to 2.0 Megabits/s upstream. Some requirements such as the carriage of video and multimedia traffic have complicated the standardisation process. A number of proprietary schemes are currently being deployed, both in the US and in Australia. Many of these have inferior performance to the proposed 802.14 protocol.

The asymmetric nature of cable modem network access is well adapted to web browsing activity, but is far less suitable for file serving, file transfer and X protocol traffic. Where a strong requirement exists for upstream bandwidth, cable modem schemes are hardly ideal.

In summary, cable modems are most often limited by asymmetric architecture, may be expensive, and may also be oversubscribed. Therefore it is unclear at this time whether the cable modem can achieve its full technological potential in the user base.

3 HDSL, SDSL, ADSL and VDSL

The HDSL (High rate Digital Subscriber Line), SDSL (Symmetric Digital Subscriber Line), ADSL (Asymmetric Digital Subscriber Line) and VDSL (Very high rate Digital Subscriber Line) schemes [Kopp, 1996c; Kopp, 1996d] are outgrowths of the established ISDN model, designed to carry significantly higher bandwidth over the established twisted pair wiring infrastructure, in comparison with a standard ISDN B-channel.

None of these services are at this time available to home users. However, we can expect to see them become available, at some yet to be determined cost premium, within the next decade.

The basic rate ISDN service provides a pair of 64 kbit/s synchronous B channels, and a single 32 to 16 kbit/s LAP-D protocol link management and packet data D channel. The LAP-D protocol is closely related in format to the X.25 LAP-B protocol. This service is termed a 2B+D.

The HDSL (High data rate Digital Subscriber Line) protocol is at this time used most often to carry 1.544 Mbit/s T.1 or 2.048 Mbit/s E.1 traffic over twisted pair cables, which are generally cheaper than coaxial and optical cables. HDSL signalling will carry a T.1/E.1 data stream in a bandwidth of between 80 kHz and 240 kHz, over cable runs of up to nominally 4,000 metres. HDSL at this time is used primarily to carry traffic between telephone switches, but may become available to customer premises in the future. The SDSL protocol is closely related to the HDSL protocol, but provides a symmetrical service for T.1/E.1 with concurrent POTS operation, over distances up to 3,300 metres. Like HDSL, SDSL is used primarily inside a communications network.

The ADSL (Asymmetric Digital Subscriber Line) is, unlike HDSL and SDSL, specifically designed to deliver traffic to customer premises. Arguably the ADSL service makes virtue of a necessity, in its asymmetric bandwidth utilisation.

Because crosstalk and interference in cable bundles is more difficult to deal with at the telephone switch end of the link (i.e. near end receiver crosstalk), the ADSL service provides much lower upstream bandwidth, in comparison with available downstream bandwidth. The ADSL model will flexibly trade bandwidth for distance. Nominal performance is to provide a downstream T.1 service to 6,000 metres, an E.1 service to 5,000 metres, a DS2 6.312 Mbit/s service to 4,000 metres, and an 8.448 Mbit/s service to 3,000 metres. Nominal upstream rates lie between 64 kbit/s and 640 kbit/s. An analogue POTS service may share the twisted pair with an ADSL service, operating concurrently.

While the ADSL service is promising, particularly in terms of downstream bandwidth to subscribers who live close to switches, it will most likely require complex and expensive modems to provide the required Quality of Service (QoS) on established wiring infrastructure, and to accommodate the complex protocol.

The ADSL service is to be followed by the VDSL service, which is at this time in definition. VDSL will fit under ADSL, in terms of achievable distance, and above ADSL in terms of achievable bandwidth. VDSL will be asymmetric like ADSL, and a very short distances downstream is expected to provide about one third of the bandwidth of a 155 Mbit/s ATM channel.

Both ADSL and VDSL are at this time immature, and if the Australian consumer’s experience with the ISDN service is a good indicator, then both ADSL and VDSL will be late to deploy in Australia, and will be most likely too expensive for wider home use. Like Cable Modem services both ADSL and VDSL are asymmetric and thus best suited for activity such as web browsing.

4 IP and the Home User

Access to the Internet requires that a user’s machine communicate with other devices using the "TCP/IP" protocol suite, comprising the Internet Protocol (IP), the Transmission Control Protocol (TCP), the User Datagram Protocol (UDP) and a wide range of other supporting protocols (eg ICMP, RIP, TFTP, bootp).

At this time the IP version in use is IPv4, which will be replaced in the next few years with IPv6, which is primarily designed to provide a much larger address space, and a range of other improvements based on extensive in-service experience gained over the last decade.

Users who employ serial modems over a POTS channel will most frequently employ the Point-to-Point Protocol (PPP), which is a HDLC based encapsulation scheme for IP, designed to run over asynchronous and synchronous serial lines. PPP may be implemented in a device driver, but is also frequently, in Unix based implementations, implemented as a process which communicates with the network protocol stack via a pseudo-device driver.

Less frequent is the use of the older SLIP (Serial Line IP) "non-protocol", which was a scheme for carrying IP datagrams over serial lines.

The PPP protocol is well suited for this purpose, and employs a Link Control protocol (LCP) to manage link state, and negotiate link configuration at the time that the link is established. Most PPP implementations also support Van Jacobsen's header compression scheme, which can very usefully reduce the traffic volume when carrying interactive terminal traffic.

Should the user have access to ISDN B channels, then a range of alternatives exist for carrying traffic across the link. The cleanest and most efficient technique is to use an ISDN interface on the machine, which will allow for synchronous transmission of data at the full 64 kilobit/s rate. Alternately, an ISDN Network Terminal can be used, providing either synchronous or asynchronous serial interfaces conforming to ITU V.110 and V.120 respectively [ITUV110, ITUV120].

In all of these instances, the communication channel between the user's machine and the dial-in site is a rather basic bit pipe. The PPP or SLIP protocol is run on the user's machine, and the router or host which it connects to at the central site. Should the user's facility include several hosts, then a common implementation is for the host connected to the network to act as an IP router.

The cable modem is a more complex device in comparison with a POTS voiceband modem or V.110/120 modem. A typical cable modem is a standalone device, which at the cable interface provides an RF front end receiver and transmitter pair, modulators and demodulators for the cable modulation, and a protocol engine with appropriate serial interfaces to recieve and transmit packets in the required format for the cable channel (eg. 802.14 packet format).

The host interface will typically be a 10 Base 2 or 10 Base T Ethernet interface, and some modems will provide a 100 Base T interface at 100 Megabits/s.

A CPU will be used to manage the state of the cable modem, and to perform IP routing between the cable and host interfaces. Typically, the management and routing processor will provide a remote login facility for the purpose of configuration and management. An SNMP remote management facility is often included.

The significant complexity of the cable modem, which is in effect a two interface router with an RF interface, suggests that it will always be more expensive than the established voiceband POTS modem and the ISDN V.110/120 NT and expectations that it can be produced as cheaply as such low speed devices are arguably very optimistic.

An ADSL Network Terminal (NT) is at this time a notional entity in the consumer marketplace. Like the cable modem, it will be a more complex and thus inherently more expensive device than a voiceband POTS modem. It will require a high speed interface to the twisted pair cable, capable of generating and recieving the complex line modulation in an electrically hostile and noisy environment. It will require a complex internal engine to handle the ADSL protocol and manage the state of the link.

The interface to the host could be via an ATM channel, or via a more conventional Ethernet interface as is done with the cable modem. Should the ADSL

NT be intended for computer use alone, with no provisions for video or other planned services, then it is most likely that such an NT would be built like a cable modem to perform as an IP router.

Both the cable modem and the ADSL/VDSL Network Terminal will provide superior performance to the POTS voiceband modem, with a significant penalty in complexity and thus cost. This cost will most likely be reflected in higher tariffs and charges in comparison with the POTS service.

It follows therefore that a service which can provide a multiple Megabit/s bandwidth at a competitive cost to a home user market would offer considerable commercial potential as it would bypass the costly fixed infrastucture schemes and defacto tariff monopolies of the CATV and telecommunications vendors.

One possible alternative to the cabled infrastructure is the use of RF wireless networking.

5 Wireless Networking Technology

Robust wireless computer networking technology is a recent development. Historically the first example of wireless networking was the Aloha network built by the University of Hawai, which utilised narrowband carrier wave techniques and in many respects demonstrated the collision detection/collision sensing paradigm which was later to become the basis of the widely used Ethernet/802.3 protocol.

Modern wireless local area networking technology is substantially more complex. It relies very heavily on spread spectrum techniques, historically a technology more closely associated with secure military communications, GPS satellite navigation and radars for stealth bombers. Wireless LANs utilise either Frequency Hopping (FH) or Direct Spreading (DS) techniques, transmitting low power signals in the 900 MHz, 2.5 GHz or 5 GHz medical/scientific RF bands. In most OECD countries, including Australia, there is no requirement to licence a transmitter in these bands, providing that it complies with the power and spectral restrictions imposed by regulators.

At this time a number of vendors supply such equipment. Typical examples are Proxim’s RangeLAN, Lucent’s WaveLAN, Windata’s FreePort and Aerocomm’s GoPrint. Throughput performance varies with product, typical figures falling between hundreds of kilobits/s and 5.7 Megabits/s [Geier, 1996]. The best throughput performance is typically achieved in the 5 GHz band, where a 125 MHz bandwidth slice is typically available for such use. Range performance is typically of the order of hundreds of metres, subject to the propagation environment and the interference environment. Equipment operating in the 2.4 GHz band is known to be susceptible to narrowband interference produced by household microwave ovens. From a cost perspective, RF LAN I/O adaptors are typically sold at hundreds of dollars per unit adaptor, which is comparable to early low cost Ethernet adaptors.

The IEEE is currently defining the 802.11 standard for Wireless Ethernet, which is expected to utilise many of the design characteristics of these existing proprietary schemes. Equipment conforming to the new 802.11 standard is not expected to interoperate with existing proprietary equipment.

Equipment for RF wireless LANs is typically built to extend an existing cable based Ethernet into building or site areas not covered adequately by the existing wiring. Wireless LAN I/O adaptors for Personal Computers are typically build as ISA, EISA or PCI bus interface cards, for insertion into such hosts. Device driver support is available in most instances for Microsoft Windows, but some cards are now also supported with Unix drivers for BSDI, FreeBSD or Linux.

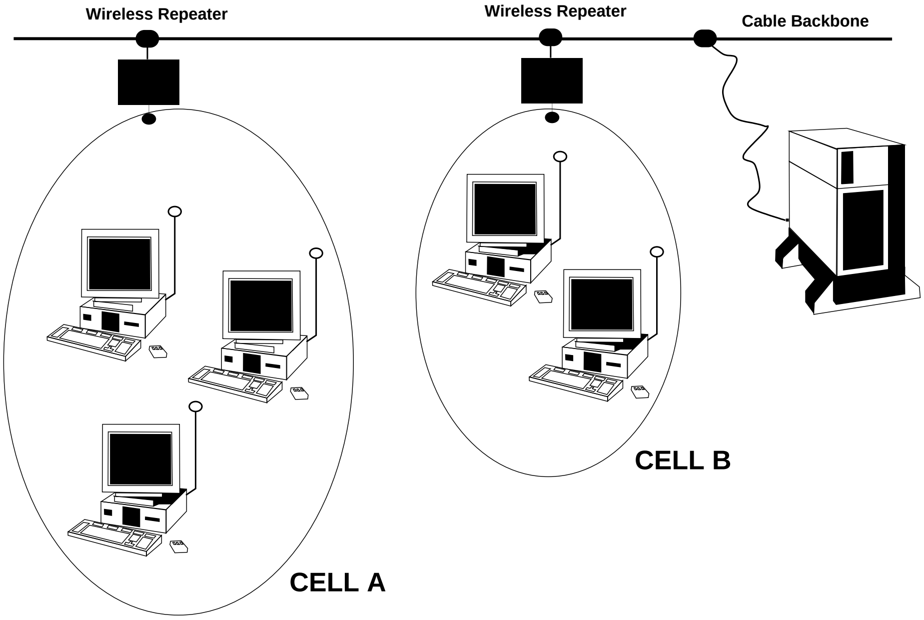

Figure 2. Wireless LAN to Cable LAN Connectivity

Connectivity to the fixed LAN infrastructure is provided by the use of RF regenerative repeaters, analogous in concept to the 10 Base T Ethernet hub or repeater (see Figure 2.). In this manner, hosts equipped with Wireless LAN I/O adapters can communicate with one another, or with hosts or servers tied to the fixed LAN.

This technology is also available for directional RF point to point links, using directional microwave antennas. Ranges of kilometres can be achieved in this fashion.

From a performance perspective, the throughput limit of Megabits/s over a symmetrical channel, is adequate for most of the applications of interest to a home user. While such performance is inferior to the theoretically achievable performance limits of a 10 or 100 Mbit/s Ethernet, in practice this nominal disparity may not be such a major issue. This is because the network I/O performance of many smaller hosts is primarily determined by the rate at which the host can generate or receive network packets. Unpublished measurements conducted by one of the authors some years ago using Personal Computers and Unix Workstations indicated that many hosts could not saturate a 10 Megabit/s Ethernet with packets, even when running synthetic throughput benchmarks such as TTCP.

Moreover, since spread spectrum techniques offer channel separation through the use of orthogonal PN codes, the available bandwidth could be utilised by code division multiplexing techniques.

Another established technology which is of some interest in this context is narrowband Packet Radio RF WAN equipment [Geier, 1996]. Packet Radio schemes such as the ARDIS network in the US, or the Ericsson developed Mobitex network utilised by US service provider RAM Mobile Data, are based upon the idea of cooperative routing. In these schemes, packets are relayed between nodes to provide connectivity beyond the direct line of site. A distributed routing protocol is used to dynamically update the routing tables in each node.

Throughput performance in these schemes varies between kilobits/s and tens of kilobits/s, and is primarily constrained by the narrowband modulation techniques in use.

Another RF WAN scheme is noteworthy of discussion in this survey [Geier, 1996, METRICOM97]. It is Metricom's Ricochet network, deployed currently in the San Francisco, Seattle and Washington, D.C., areas of the US. The Metricom network uses 900 MHz band power pole or rooftop mounted repeaters, which use a packet radio style distributed routing protocol. Users with Metricom RF modems are provided with a 28.8 kilobits/s serial interface to run PPP or SLIP to remote hosts. Connecting directly between a pair of modems, with no intermediate hops, throughputs of about 40 kilobits/s can be achieved [METRICOM97]. The Metricom networks employ central network IP routers, and users are dynamically assigned IP addresses with every connection, from a pool of available router ports. In effect, the Metricom network emulates a telephone dial-in connection, using packet radio to replace telephone wiring.

It is clearly evident that the ideas embodied in RF Wireless LANs and Packet Radio schemes offer further potential for exploitation. The Suburban Area Network (SAN) model aims to achieve this.

6 The Suburban Area Network

Conceptually, the Suburban Area Network (SAN) is a distributed RF network which exploits packet radio cooperative routing over Megabit/s spread spectrum microwave communications channels. Households participating in an SAN would maintain a roof mounted omnidirectional antenna which would provide connectivity to peer stations within achievable radio range for the propagation environment in question. It is also possible to use high-gain, more directional antennae to improve the range, however the spectrum management authorities limit the RF power density for unlicenced operation.

A SAN would provide a suburb wide cell within which all hosts can communicate with one another, either directly, or with several hops, allowing a throughput of a similar order of magnitude to that of a direct spread spectrum connection. Should existing spread spectrum LAN technology be exploited for this purpose, symmetrical throughput of the order of Megabits/s will be achievable.

The SAN would connect to the cabled network infrastructure through one or more routing nodes, which would have a connection to the cabled network. Ideally, such a connection would be at several Megabits/s, such as a T.1, E.1 or Ethernet connection, or a wireless datalink connection. Should the number of hops be small between a station and the routing node, and should aggregate traffic levels be modest, then the station can have a symmetrical connection to the cabled network with a throughput between hundreds of kilobits/s and Megabits/s.

The SAN would thus offer throughput performance which is similar to that provided by cable modems and ADSL connections, far superior to that provided by voiceband POTS modems and ISDN-B, with the important attribute of providing symmetry in downstream and upstream channel speeds. Therefore the SAN would be a more suitable network for general purpose use, as compared to the web browsing specialisation of the cable modem and the ADSL/VDSL network. The SAN is inherently more suitable for applications such as employees working from home, or for the provision of networking services to academic staff and students. Symmetry will allow for the use of protocols such as NFS, CIFS, IPX/SPX and X11 which are of central importance should we seek to provide the full utility of the network to the home user.

The issue of heavily asymmetric user loads produced by web browsing must be addressed, as it has the potential to degrade the performance of any SAN cell. A caching HTTP proxy server at the interface between an SAN cell and the network could be employed to alleviate but not eliminate this potential problem.

An important attribute of the SAN is that it decouples the home user from the cable television and telephone wiring infrastructure, thereby bypassing the costs associated with renting these services from large and thus arguably uncompetitive providers. Should we assume the use of existing RF LAN technology, the cost of a single user installation will be of the order of $1000, comprising the cost of the interface card or modem, household cabling and antenna. The only recurring cost will the the throughput charges associated with the shared high speed connection between the SAN cell and the fixed network. Should this connection be provided by the employer or academic institution sponsoring the network, using for instance a point-to-point spread spectrum microwave link, then the cost of the connection comprises mainly the non-recurring cost of the microwave link.

Even should this service be provided by an Internet Service Provider (ISP) rather than employer or academic institution, the costs will be highly competitive. This is because the cost of a high speed cable connection and its maintenance are not incurred.

A number of important technical and design issues must be addressed in the development of the SAN model:

- choice of RF modulation scheme, examples being FH and DS, and whether to use off-the-shelf 802.11 wireless Ethernet or whether to use a unique protocol for host to host connections.
- choice of PN codes, management of PN code space, and security of traffic.
- choice of routing scheme and protocol. Because the SAN is in many respects a static network, existing IP level packet routing may be directly exploitable. Otherwise, packet routing at the RF datalink layer may be required.
- network management and troubleshooting. Because the individual nodes are owned and operated by individual users, arrangements must be made to accommodate this.
- network robustness, as topology under some circumstances will produce situations where a particular node may be the cut vertex of the connectivity graph for the SAN.
- RF propagation performance under adverse weather conditions. Heavy rain and fog can significantly attenuate microwave transmission [NASA1108] and this may impact QoS in some geographic environments.
- resilience to narrowband interference. Leakage from microwave ovens and other microwave sources could impair network performance, particularly where the topology is unfavourable.
- should the SAN node be implemented as a plug in I/O card for a home computer, or as a separate SAN RF modem with a discrete networking interface, as is done with many cable modems ?

It is evident that a wide range of technical alternatives exist for the design of an SAN, offering a broad range of performance and QoS alternatives.

6.1 Suburban Area Network Node

There is a diverse range of possible implementations for a SAN Node, which is the basic building block in any SAN. Conceptually, the SAN node is a device which can perform the functions of cooperative routing, as well as providing datalink level radio or other connectivity to peer or remote nodes.

The simplest implementation for an SAN Node is to fit suitable communications interface cards into a personal computer, and provide appropriate protocol stack and device driver software to implement the functions of the SAN Node. This scheme clearly incurs the lowest possible hardware implementation cost. This implementation strategy however has a number of important weaknesses:

The host platform incurs the load associated with the routing of traffic and processing of the SAN protocol stack which will necessarily impair the performance of other applications running on this platform.

Because an arbitrary SAN Node may carry traffic for a number of its peers, it is imperative that the host platform deliver sufficient performance to handle the SAN traffic load. A suitable multitasking operating system kernel is essential, to ensure that the processing of routed traffic is prioritised above the processing of local tasks.

The reliability and availability of the operating system are of vital importance in that a system crash has the potential to affect the operation of other nodes in the SAN. A well designed SAN with many redundant internal paths will have a good level of resilience to such node failures, but will suffer a loss in aggregate throughput as a result.

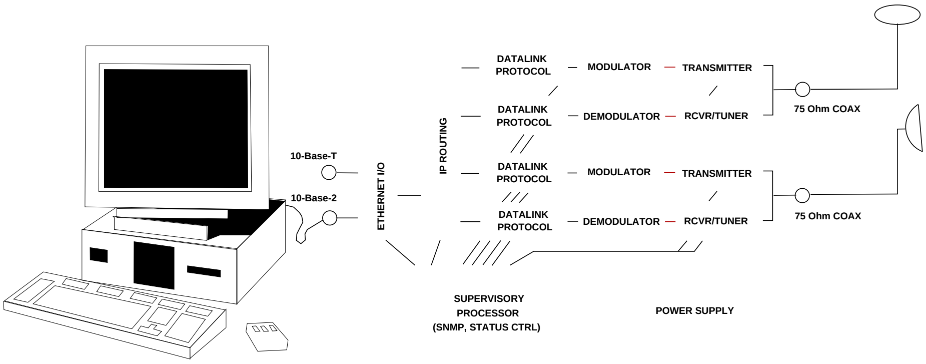

Figure 3. Standalone SAN Node

A standalone SAN Node (see Figure 3.) implementation addresses many of these issues. Such a device will appear to a host much like a cable modem. Functionally the standalone SAN node appears more like a wireless RF LAN repeater .

Our proposed implementation includes two fully redundant RF links, an embedded microprocessor based protocol engine and router, which includes SNMP management support. We envisage a device similar in size and power consumption to an RF modem, capable of unsupervised continuous standalone operation. This decouples the SAN from the local computer, allowing the latter to be shut down at arbitrary times.

It should be noted that each participant in the SAN owns and controls his own SAN node. The success of the SAN paradigm is based on the cooperation of the SAN participants, but does not rely upon any individual participant.

7 SAN Implementation Issues

The Suburban Area Network is beginning its life as a Monash University research project. There are many unresolved questions and problems. While spread-spectrum microwave technology has been identified as suitable for this application at the physical link layer, it remains to be seen through experiment, what range can be achieved with acceptable error rate and still meeting licencing restrictions.

While the data links provide good engineering challenges, clearly the more difficult aspects of the project are those related to network topology and to the routing protocols. Here we want to achieve a good quality of service but want to minimise the cost. Routing will have to be fairly dynamic since we will not have great confidence in the reliability of any single node, given that individual nodes are located in people's houses and subject to the usual vagaries of power supply and accidental switching off. Thus ideally one would not want to be reliant upon a single routing node. If the nodes are situated such that communication to multiple nodes is possible with a single omnidirectional antenna, then it is relatively easy to provide suitable routing protocols to route around failed nodes. However if one is pushing for range then it is not sensible to provide too many directional antennae pointing to other nodes. The difficulty then is to provide a routing system which maximises the quality of service with the constraints of limited connectivity. Much of the research in this project concerns this problem.

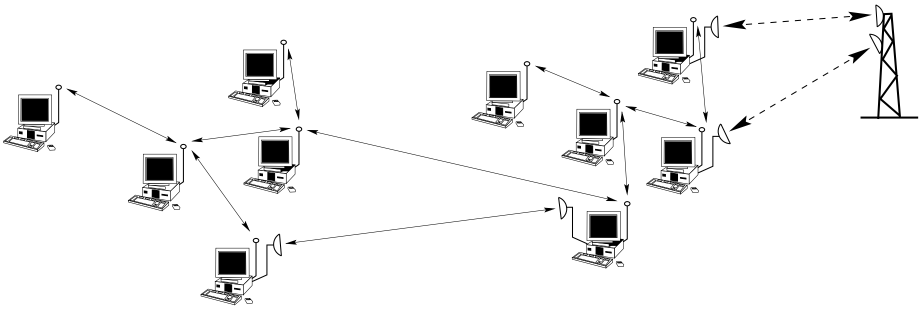

Figure 4. Suburban Area Network

We anticipate that a SAN Node implementation with two RF links is both practical and economically feasible. Such a node will provide sufficient redundancy in connectivity to allow for load balancing between nodes and allows the maintenance of connectivity in the face of any single link failure.

8 Conclusions

This paper has surveyed existing technology for the provision of Internet connectivity to the home user, and for the provision of RF wireless networking. The limitations of these schemes in throughput performance and utility have been discussed.

The Suburban Area Network (SAN), a scheme using cooperative packet routing and high-speed spread spectrum radio links, has been proposed as an alternative to existing techniques for the delivery of Internet services to home users.

The SAN model offers the potential for multiple Megabits/s symmetrical throughput performance, bypassing the fixed wiring infrastructure and incurring modest recurring operating costs.

Further research will be required in the area of link transmission protocols, cooperative routing protocols, robustness of link performance and topology, and implementation for minimum cost.

References

Flock W.L.(1987) Propagation Effects on Satellite Systems at Frequencies below 10 GHz, Second Edition, NASA reference Publication 1108 (02), 1987

Geier J. (1996) Wireless Networking Handbook, New Rider’s Publishing,Indianpolis, 1996.

ITUV110 - [V.110] Recommendation V.110 (10/96) - Support by an ISDN of data terminal equipments with V-Series type interfaces

ITUV120 - [V.120] Recommendation V.120 (10/96) - Support by an ISDN of data terminal equipment with V-series type interfaces with provision for statistical multiplexing

Kopp C. (1996a), An Introduction to Cable Modems, Open Systems Review, Sydney, August, 1996

Kopp C. (1996b), The Anatomy of a Cable Modem, Open Systems Review, Sydney, August, 1996

Kopp C. (1996c), Twisted Pair Fights Back - Advanced Digital Subscriber Line Technologies, Open Systems Review, Sydney , September, 1996

Kopp C. (1996d), Digital Subscriber Line Technologies - The Technical Issues, Open Systems Review, Sydney , September, 1996

METRICOM97 - Metricom Web Page at http://www.ricochet.net/

RFC1055 - Romkey J.L., Nonstandard for Transmission of IP datagrams over Serial Lines: SLIP, Request For Comment, Internet Engineering Task Force, June, 1988

RFC1661 - Simpson W., Editor, The Point-to-Point Protocol (PPP), Request For Comment, Internet Engineering Task Force, July, 1994

RFC1094 -Sun Microsystems, Inc, NFS: Network File System Protocol Specification, Request For Comment, Internet Engineering Task Force, March, 1989

RFC1813 - Callaghan B., et al, NFS Version 3 Protocol Specification, Request For Comment, Internet Engineering Task Force, June, 1995