Issues in Residential Broadband Internet Service Provision

Saleem N. Bhatti <S.Bhatti@cs.ucl.ac.uk>

Graham Knight <G.Knight@cs.ucl.ac.uk>

Abstract

The CATV industry sees itself as extending to become a FSN (full service network) provider, and so would like to offer data services, notably IP (Internet Protocol) connectivity. Cable companies have the raw network capacity to the home that gives them the potential to offer broadband services that can not be matched by traditional PSTN/modem connectivity. In the home, SoHo (small-office/home-office) and SME (small-to-medium enterprise) environments, use of applications such as electronic mail and the World Wide Web (WWW, the Web) are proving to be very popular and effective. Also applications allowing such functions as conferencing are becoming available. Broadband residential services will enable cheap accessibility of such functions from the CPN (customer premises network). Internet users have the following requirements:

- a familiar integration interface for the CPN
- a flexible service offering
- secure communication
- addressing and routing should follow current and evolving Internet models
- full Internet services

This paper highlights some of the major issues involved in considering the provision of such residential broadband Internet services, notably:

- a suitable network infrastructure
- provisioning of connectivity
- availability of technology

These issues are being investigated by the authors as part of the European Framework IV ACTS project IBCoBN (AC101).

I. Introduction

Most current home, SoHo (small office/home office) and SME (small to medium enterprise) users (we will group them as residential users in this paper) must use a modem and a normal analogue PSTN line to connect to the Internet, via their ISP (Internet service provider). Affordable computing equipment, high competition amongst ISPs and cheap PSTN access via local PoPs (points of presence) have made such connectivity easily accessible to the population at large. In many cases, this has resulted in poor performance of servers, slow access speeds to some providers and over-subscription. Two of the these factors can be overcome by sufficient resourcing by the ISP[1] – the one that can not be overcome is the access speed to the ISPs. The PSTN analogue local loop is becoming saturated and it is unlikely that modem speeds can increase significantly[2].

ISDN (integrated services digital network) has had limited implementation, with some countries providing greater coverage that others [1]. Leased-line solutions are generally too expensive for home users and for many SME users. The summary in Table 1 shows the typical snapshot of costs of these services in the UK at the time of writing.

Service	Data rate [Kb/s]	Installation costs	Fixed running costs (per year)	Call charges (per minute)[a]	Interface or adapter cost
PSTN	28.8[b]	175.00[c]	150.00[d]	0.02	120.00
BR-ISDN	64/128	650.00[e]	600.00[e]	0.02	400.00+[f]
Leased-line[g]	64+	1500.00+	6500.00+	–	1500.00+

Table 1: Typical connectivity service costs in the UK (all costs approximate, in ECU)

Notes for Table 1:

a) Local rates, off peak – nearly all ISPs offer local PoPs. Some CATV networks provide free local calls during off-peak hours.
b) Capability of highest speed modems using proprietary technology is actually 33.6Kb/s but the current standard is 28.8Kb/s (using compression).
c) This is the cost if you need to have a phone line installed – most people do have a phone line and so one might consider this cost to be zero.
d) BT residential customer rate.

[1] Experience in the US has shown that free local phone calls lead to lengthy Internet access, which has lead to heavily congested services.

[2] 1997 sees the arrival of "modems" with data receive rates of 56Kb/s and transmit rates of 28.8Kb/s. However, such modems will require new digital tranceiver equipment to be present at the ISP and one could argue that these devices are not strictly "analogue modems".

e) BT "Low-Start" rate.
f) ISDN adapters vary in price from passive adapter cards that just offer bit pipes, to sophisticated ISDN TAs (terminal adapters). This price is for an adapter card giving fairly basic connectivity.
g) This is the cost for 64Kb/s. Rates up to 2Mb/s are available but the cost is considerably more.

Not all ISPs have ISDN or leased-line access. There is clearly a market for a low cost (of the order of the PSTN service) high capacity data service for the residential market. The cable companies have the raw network capacity in the ground to provide it. If a suitable network architecture can be found, coupled with a suitable pricing policy, CATV network operators have the potential to exploit the growing Internet services market. This could happen either by becoming ISPs themselves or offering high quality, high data rate connectivity between existing ISPs and CATV subscribers. There is the potential to provide differing degrees of connectivity (data rates and services) to residential and business users in a flexible manner.

This paper looks at some of the service and technology issues faced by the cable companies in providing data services and in particular Internet connectivity to the residential market.

- Section II lists the requirements of Internet users.
- Section III examines the issues for provisioning Internet services in a residential environment.
- Section IV looks at the technology constraints that apply, including the CPN, the provider network and IP itself.
- Section V provides a short summary of the discussion in this document.

II. The requirements of Internet users

Typically, the services required by today's Internet users with high-capacity access are:

U1. real-time, interactive communication: many of the services that today's users would like to have involve them being able to interact closely and in real-time with that service (e.g. audio/video conference).

U2. access to multimedia resources: users would like to have network access to sound, still pictures, video images, graphics and information databases as easily as they can listen to their stereo, watch a video recording, or use a multimedia CD-ROM.

Additionally, in the short-to-medium term, the following will also be necessary:

U3. secure communication: users may wish to protect their transmitted data from being observed, modified, forged or replayed by a hostile third party, and have similar assurances about the data they receive. Parents may wish to impose controls on the sites that their children can receive data from.

U4. many-to-many (group) communication: people need to interact with each other in order to work (and for social reasons), so the network should allow multi-party communication.

U5. flexibility in use of the network: users do not wish to pay for a high capacity service which they rarely use, however, they would like the ability to have high data rates at times, on demand.

From this list of user requirements, we can write the following functional requirements:

F1. high speed, low latency communications: to support real-time, interactive applications. (Addresses U1 and U2.)

F2. a data-transparent, integrated services network: the network architecture should allow the transport of many types of data on the same network. (Addresses U1 and U2.)

F3. security services and mechanisms: deployment of facilities such as encryption, digital signatures and exchange of certified tokens to give privacy, proof of origin of data, and authenticity of sender. Configurable firewall technology for filtering data would also be useful. (Addresses U3).

F4. multicast capability: to allow easy group communication. (Addresses U4.)

F5. resource reservation mechanisms: IP users (applications) should be able to signal the network of their resource/traffic and QoS requirements for a particular media flow or logical (IP based) connection. (Addresses U5.)

A similar list is also presented in [2]. Note that the provision of certain aspects of F2 may rely on the provision of F4 and F5. The provision of F5 would be useful to the network provider for the purposes of network management. It could also be used for accounting and billing purposes. Mechanisms for providing F3, F4 and F5 inevitably require some state information to be maintained somewhere in the network. Scaling issues concerning such information have to be considered carefully – the amount of state information that is required, how it is distributed and how often it is refreshed. Discussion of such issues is outside the scope of this paper.

III. Residential service requirements

It is instructive to look at the current model of some typical residential users:

- Today's home user will probably be an Internet enthusiast who has an interest in computers. Such a user may only require F1 and F2, although parents may wish to use some sort of security service to restrict access to their children (F3). Many such users will be hobbyists and will not have requirements for the functionality offered under F3, F4 and F5.
- The SoHo user would make use of (multimedia) e-mail services and file transfer, and possibly be able to access some technical information available somewhere within the Internet. If this user is a teleworker, his/her employer may make resources available for use electronically via file servers and information servers such as FTP, Gopher and WWW.
- A SME user might have a leased line connection to an office and provide internal IP-based services (Intranet services) as well as offering public access to some services such as WWW and FTP servers. Such a user may support access to SoHo workers and teleworkers.

SoHo and SME users are business users, and they may need to use functions offered under F3, F4 and F5 for the purposes of their business. A home user may not need these functions. This distinction is important for the provider in determining the way in which services should be deployed and offered to potential users. Many home users may comment that, "My 28.8Kb/s modem is too slow for my needs." This may be true, but it is also likely that the Internet resources that they are trying to access with such a modem are in a (geographically[3]) remote location, to which that user has slow access. In many cases, the bottleneck is not the local-loop but is somewhere within the "Internet cloud". The SoHo and SME user, however, is likely to be relatively close to the resources with which they interact and the people for whom they provide services, e.g. the teleworker scenario.

Additionally, in order to appeal to a wider range of home users, and not just the Internet enthusiasts, the Internet technology needs to be accessible to a much wider range of people. For example:

- access to local services and information, e.g. local council information, library holdings etc.
- communication between local businesses, e.g. some towns have internal mail services for document delivery. This could be achieved electronically.
- local transport information timetables, travel conditions, etc.
- home banking, home shopping, entertainment schedules, etc.

These services are **neighbourhood services**, and could be based around information held very locally. However, the kinds of people who need these services often do not have access to the equipment required to obtain these services. An interface that uses a normal PC, keyboard and mouse is too complex for users accustomed to a remote-control handset. This could change in the near future. If this kind of network technology proves to be the best delivery system then suitable interfaces will develop.

III.A Neighbourhood Intranet

The discussion above outlines a current scenario that can be summarised in the following points:

N1. **High speed connectivity is not often of great benefit to home Internet users:** such users access remote services where the data has to traverse paths which include links that are slower than the local link to the PoP. The many national and European initiatives which encourage the deployment of fast data networks may change this in the future, but in the short- to medium-term, much of the wide-area Internet remains slow and lossy.

N2. **SoHo and SME users often have need for high data rates in a localised area:** teleworkers businesses and organisations that offer local-area and metropolitan-area services could benefit from the services and information offered via local and metropolitan high-speed networks.

N3. **Provision of local information and services:** these could be provided and accessed at high speed if the data was routed locally, i.e. within a neighbourhood space.

N4. **Residential users need access to infrastructure:** residential users do not have the correct equipment and this could still be true even if N3 could be put in place.

Cable TV companies in Europe are considering the provision of data services as they aim for their networks to become FSNs. Some Cable TV companies in the USA are already offering high speed data services based on a neighbourhood LAN model. They offer shared-media access to data rates from several hundred Kb/s to about 10Mb/s. One way in which such networks could be used is by addressing the points N1 to N4:

[3] Geographic location is not necessarily indicative of ease of access via IP, but can often be taken as a good hint – the UK to US link, for example, can sometimes be a few hundred b/s so access to US sites is generally slow. However, it is possible that a local PoP number for an ISP is routed for hundreds of miles before accessing local resources.

- **Allow a range of connectivity services:** caters for the home user, the SoHo user and the SME with a different set of connectivity and service packages.
- **Encourage local businesses and organisations:** the available network capacity is best used locally, so the services and information should be local too.
- **Provision of technology to allow access to local services:** provide a box in the user's home that is easy to use but allows access to the local information and services (e.g. "network computers" operated from a simple handset).

The ACTS project IBCoBN (AC101) involves European cable companies and equipment manufacturers. They are looking at the provision of a data channel on the existing (and future) CATV infrastructure, in order to provide broadband services using ATM (asynchronous transfer mode). To carry IP over ATM is known technology, and so it becomes possible to offer broadband IP services locally. One scenario of such an offering is to consider a **neighbourhood Intranet**. This would have its own local servers (e.g. FTP, WWW) and offer high speed, low cost access to this information.

III.B Cost

There are still some factors concerning cost that need to be resolved. While it would be expected that the CATV operator would pay for the network installation and devise a suitable charging structure, many subscribers might still be reluctant to pay for an advanced STB and home computer themselves. While the costs of home computers are decreasing, it can still cost around 1000ECU for low-end multimedia PC and suitable software. This would discourage the widespread use of a neighbourhood Intranet. If the cost was met by some sort of (local, national or European?) funding initiative to provide such equipment, then perhaps the take-up of these services would be more successful, e.g. akin to the MiniTel deployment model in France. Another model might be to try and encourage the cable operator or service provider to provide the user's access equipment free of charge and recoup the cost through the charges for services. However, current trends seem to indicate that "network computers", which can use the TV as a display and offer a browser-type interface, may be available for around 400ECU and the end user would be expected to pay this cost.

IV. Technology issues

The previous two sections looked at some requirements and service issues. Here, we discuss issues concerning the technology required for broadband IP services. In particular, we focus on some specific questions in connectivity technology (for both the CPN and the provider network) and the use of IP.

The CATV community is considering the use of ATM in the sub-network, and is especially looking to offer digital entertainment channels in the future. ATM has been designed to be data transparent, and allow the use of QoS control features. The mainly analogue networks services, will in the future, be delivered in digital form. This will allow very flexible use of the network, as well as making it possible to offer services such as VoD (video on demand). The transition from analogue to digital will not result in the immediate use of baseband ATM transmission: the existing analogue channels will be used, but will carry ATM streams. The use of ATM in the provider network raises some issues that are addressed in the rest of this section.

IV.A The CPN and the access network

The CPN (customer premises network) may be just a single computer, in the case of the home user, or a LIS (logical IP sub-network), in the case of the SoHo and SME user. The manner in which the CPN connects to the provider network may have implications for the way in which the routing and addressing is achieved within the CPN. For example, if the connection is a traditional "dial-up" service, the addressing and routing architecture within the CPN may be different compared to when the interface to the CPN represents a direct Internet connection.

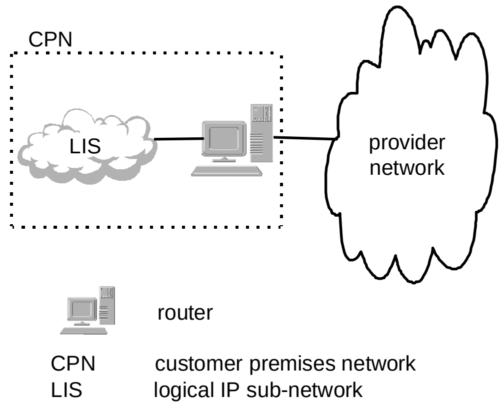

Figure 1 Connection of a logical IP subnetwork

A SoHo or SME user, may want to connect a LAN, such as Ethernet, to the provider network. This would make the CPN a LIS, and so we need a routing function between provider network and the CPN, as recommended in RFC1009 [3] (Figure 1). Indeed, many companies use such a configuration as an IP firewall. The routing function could reside within the customer's equipment. However, if it is realised within the STB, this would allow the cable operator to have important management control at the IP level, as most cable companies currently own the STB.[4]

From the CPN to the provider network, it makes sense to adopt an interface that is currently used in (or is easily adaptable to) a LAN environment, in order to allow ease of connection to prospective users. The use of a MAC interface would allow existing IP applications to run unmodified. However, for the future provision of F4 and F5 (and possibly F3), the use of this interface may be need to be modified.

To allow the CPN to be connected to the cable network, we require some equipment that will perform a similar function to the one which a (normal) modem performs for the PSTN – transform and modulate the digital information for transmission on the cable and also decode data from a similarly received signal. This requires a RF (radio-frequency) modem, commonly called a **cable modem**. There are two main points of interest concerning cable modems:

a) physical/electrical interface/protocol (level 1)
b) link level interface/protocol to use (level 2)

Neither of these functions are currently documented as ratified standards within the cable community. However there is much activity to resolve this problem under the charge of various standards bodies. The main players here are the cable companies themselves, setting up their own proprietary solutions with a wide variety of service offerings, with data rates ranging from 128Kb/s to 10Mb/s on a shared media. These follow a neighbourhood LAN model.

IV.B The provider network

It is useful to consider some of the issues within the set-up of the provider network, because they directly affect the CPN. The first consideration is that we have assumed that the provider network will use a "suitable" sub-network technology, however we have not yet determined what is "suitable". The data services must be built bottom-up on existing network infrastructure.

Data access to the CATV networks will be via a RF modem that operates in the region 5MHz to 40MHz (approximately). There exists some bandwidth in the range 550MHz to 850MHz (approximately), but current plans are to use this for digital TV channels. However, as these channels are seen at the CPN, some of these may also be usable for data (e.g. for a separate downstream data channel). The amount of bandwidth available to provide the upstream capability is fairly restricted. The use of this upstream channel would be shared by the all the users connected to a common local distribution point, and so the scaling potential of any system used is an important factor.

[4] The subject of who owns and/or pays for (and controls) the STB is considered by some to be a political issue.

IV.B.1 The ATM network model

If a MAC interface is chosen for access from the CPN to the provider network, a function that could map between the ATM interface and the MAC interface would be required, for instance the ATM Forum's LANE Emulation model [4]. The interface offered at the STB would be important for the residential user too; an ATM adapter card for a PC can currently be up to an order of magnitude greater in cost that an Ethernet adapter card. If a non-MAC interface is offered, (ideally) a transparent address translation/resolution and protocol mapping would be necessary via the STB or somewhere in the provider network.

The other option, which might be used in the short term (especially for residential users), is to offer dial-up type access, in a similar model to today's method of using a modem to connect to an ISP, but of course, at higher speed, for instance by use of an ATM SVC (switched virtual circuit). In this case, a standard interface to the SVC would be required, otherwise each CATV network operator would need to provide the software required by the subscriber's machine for each particular cable modem, as well as the cable modem itself.

For the purpose of deploying IP, the IETF (Internet Engineering Task Force) has set up several working groups. The IETF working groups (WGs) of interest are:

- **IPATM (IP over ATM) WG:** developing standards for routing and forwarding IP packets over ATM subnetworks.
- **ROLC (Routing over Large Clouds) WG:** develops standards for wide area internetworking.
- **INTESRV (Integrated services) WG:** developing standards for the end-to-end controlled QoS for integrated services over IP.
- **Network WG:** responsible for forwarding general Internet standards, including the current activity to define the new version of IP, IPv6.

Considering first F1 and F2, an excellent exposition of the issues is given in [5], which also serves to summarise some of the IETF work. The "Classical IP over ATM" model is documented in [6]. RFC1577, proposes the use an of ATM address resolution protocol (ATMARP). The LIS routers could also be used as ATMARP servers. The [6] solution has advantages in that it is a simple and widely used model, employing mechanisms that are current practise, and so the strengths and weaknesses of this set-up are well understood. This is not the only solution offered.

Another solution to the ATM connectivity problem is proposed by the ATM Forum in the form of LANE (LAN emulation). This transparently emulates a MAC-based subnetwork at Layer 2 over an ATM network. Whilst having the advantage that it presents a familiar LAN interface to attached stations, the current LANEv1 does not scale well to the wide area, and interconnection between LISs still makes use of routers in the same way as for Classical IP over ATM. These deficiencies are being addressed by the ATM Forum in the specification of LANEv2.

User-to-network signalling is supported at the ATM UNI (user-network interface) in the form of the ATM Forum's UNI specification. The signalling system for ATM uses Q.2931, whilst most Internet users employ RSVP (Resource Reservation Protocol) [7] or ST2+ [8]. There is activity within the IETF to define a method of mapping RSVP primitives to Q.2931.With ATM, therefore, it seems easily possible to provide F1 and F2, with solutions to F4 and F5 receiving much attention and expected to mature rapidly.

IV.B.2 IEEE P802.14 WG and ATM Forum Residential Broadband (RBB) WG

The IEEE P802.14 work and ATM Forum's RBB WG have been looking into the provision of broadband capability in the same domain as CATV networks. Most of the current work focuses on the provision of ATM services on HFC networks. A summary of the P802.14 work is presented in [9]. [9] also states that the two competing protocol stacks are as shown in Figure 2 (the "ATM Friendly" stack and the "ALL ATM" stack), suggesting that the ATM Friendly stack is likely to be chosen as its access interface is consistent with both the ATM Forum model and the IEEE 802 LAN model. This would also then allow easy transition of IP applications onto such networks, and allow use of the neighbourhood LAN model (see below). This will provide F1, F2 and F4 but further investigation would be required to comment on suitability to provide F5.

In Figure 2, the upper part of the diagram has been added by the authors to show how IP could be provided at the appropriate interfaces. IP would be able to work over either stack, though in the ALL ATM stack, the ATM interface may not be available as depicted.

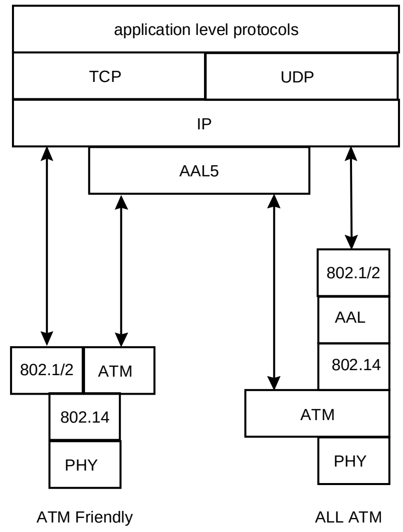

ATM asynchronous transfer mode
AAL ATM adaptaion layer
PHY physical layer
IP Internet protocol
TCP transmission control protocol
UDP user datagram protocol

Figure 2: Possible protocol stacks for provision of IP over IEEE P802.14 and ATM Forum RBB

IV.B.3 The neighbourhood LAN model

Another connectivity strategy, currently being brought into service in the US, is the use of neighbourhood LANs. As the name suggests, these offer full-time connectivity, residential broadband data services. A typical set-up is shown in Figure 3.

This set-up uses a RF modem to connect to the existing cable infrastructure in order to realise a direct IP connection. The machine directly connected to the RF modem additionally needs an adapter card. The IP traffic is taken to the Internet via a hierarchy of IP routers. The ISP which eventually takes the IP traffic need only furnish a single connection to the CATV provider network, as opposed to the many hundreds it may need if offering the IP services using a dial-up approach. The maximum coverage of such a LAN is expected to be about 200Km in round-trip length.

This model has the advantage that it is easily installed on existing infrastructure, with little or no extra switching and transmission capability required (however support for use of the return path within the network and at the head-end is required, as well as the provision of the IP-based application services, of course). Unfortunately, due to the restricted spectrum allocation for the upstream channel in some installations, the IP data must compete with interactive TV, telephony and other data services for bandwidth. The current disadvantage with this system is that there are no standards for such connectivity and the solutions are all proprietary. Work is underway within IEEE – IEEE 802.14 Forum – to define standards for such residential broadband access, but these could take 2+ years to mature.

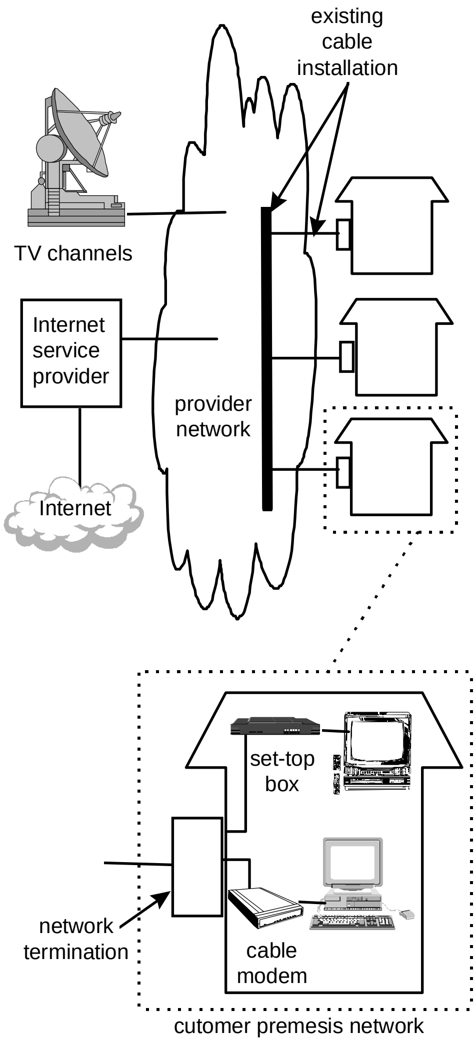

Figure 3 Neighbourhood LAN model

If LAN connectivity is available, then solutions and mechanism for providing F1, F2 and F4 for the Internet user over LAN environments are known and these could be adapted relatively easily for use over the neighbourhood LAN. However, the realisation of F5 needs some attention in the LAN environment.

IV.B.4 The DAVIC network model

DAVIC (Digital Audio-Visual Council) is also looking at the provision of Internet services within its own framework [10]. In fact, the Internet protocols – TCP, UDP and IP – form part of the DAVIC specification for use in signalling and network management. There is the possibility that the DAVIC considerations may result in the use of MPEG-2 as a level 2 mechanism. The system proposed allows the subscriber or application to select an Internet service over a DAVIC network. In this model, Internet access is available as would be any other DAVIC service. It can be individually selected, and all traffic for that service (i.e. IP packets) is directed via a server providing an IP service. Use of such a set-up, whilst conforming to the DAVIC model, is not fully investigated or well understood with respect to the functional requirements F1 to F5 listed previously. This model would probably allow F1 and F2 to be realised by defining suitable encapsulation methods. It seems plain, however, that provision of F4 and F5 could be very difficult. Although DAVIC does support U-N (user-network) signalling, this is provided by using DSM-CC (Digital Storage Media – Command and Control) [11]. To allow IP users to make use of this, the RSVP and/or ST2+ protocols (see below) would need to be mapped to DCM-CC, and mechanisms to enable such facilities are not supported by the DAVIC architecture. Furthermore, with any central server-based system, there are bound to be scaling problems. However, such a service may not be perceived by a subscriber to be very different from the current dial-up services using modems and the PSTN. This may offer a short-term solution to providing simple, but useful, IP connectivity.

IV.B.5 The IETF IPCDN model

At the time of writing, the IETF had recently set up a WG to look at the provision of IP connectivity over cable networks [12]. This WG is called IP over Cable Data Networks (IPCDN) and will examine the IP routing, addressing, address resolution and connectivity. This WG is currently establishing a network model to act as a framework in which it will define procedures for providing IP. IPCDN WG has a different objective than the other models mentioned in this section – it is looking more at level 3 (IP level) issues rather than level 1 (physical) and level 2 (data link) issues.

IV.B.6 Network management

It is proposed that SNMP (simple network management protocol) [13, 14] be used for the management of DAVIC-based networks and services. This is not an unreasonable choice considering that SNMP is considered widely to be the *de facto* industry standard for network-element management. There are already MIBs (management information bases) defined for the management of ATM systems [15, 16] and DAVIC has also produced MIB definitions [17, 18] for use on its own network model. SNMP is an Internet protocol and uses UDP (user datagram protocol). So, if SNMP is used for network management and the management functions required are extended to the STB, then there will be IP connectivity to every cable user! It is unlikely that the current IPv4 address space could cope with this demand[5], but the standards for IPv6, with its greatly expanded address space, are now agreed and IPv6 implementations and products are now appearing. Site-local and link-local IPv6 addresses could be used to great effect in a neighbourhood Intranet.

IV.C The Internet Protocols

The current Internet Protocol is IPv4 [19]. This is undergoing some considerable revision at the moment, and there are standards proposed for the new version of IP – IPv6 [20]. IPv4 has been used to provide a best-effort service. A significant change from IPv4 to IPv6 is that the current 32bit IP address space will become 128bit. Along with this is the proposal that the current flat routing system be updated with the use of the new address space to provide a hierarchical addressing mechanism that would make routing simpler, more efficient and easier to manage [21]. This may have implications for the address space which a SME might wish to use within its CPN. However, a detailed discussion is outside the scope of this paper.

Current IP technology usually provides F1, but the exact nature of the connectivity may affect the service received by the user – a UCL connection into JANET or SuperJANET (ATM-based academic network) provides F1, whilst a residential user using dial-up access would not get the same level of service. Hence, the provision of F1 from the CPN and provider network is important. The DAVIC model, ATM model and the neighbourhood LAN should all be able to provide F1. Additionally, RFC1889 [22] documents RTP – a transport protocol for real-time applications which is designed specifically to cater for such uses as interactive, real-time audio, video and simulation (e.g. virtual reality). IPv4 was designed to allow any data type to be transmitted, although it was mostly designed for robustness in the face of partial network failure rather than to provide facilities for real-time traffic. Coupled with some of the resource control and QoS control mechanisms, and the protocols mentioned below, IPv4 can provide F2 and IPv6 has been designed with such uses in mind. Additional work specifically looking into an integrated services model for the Internet, that pre-dates the IPv6 work, is documented in [23]. Further discussion on the provision of real-time services within an ATM-based are presented in [24].

IP networks have often been criticised for their lack of security (F3). Whilst some of this is directed at applications (e.g. SMTP mail), IPv4 does not have any native security features. However, IPv6 does have native security features [25] providing privacy, authenticity and data integrity mechanisms. Firewall technology is also available [26] to provide site-based packet filtering. The work in [25] and the use of [26] can also be applied to IPv4.

Multicast capability (F4) has been available on the Internet for some time. It is possible to tunnel multicast communication even over non-broadcast (point-to-point) networks. For the CPN, MAC LANs have native broadcast facilities which can be directly used to support multicast. Over the wide area, other considerations are necessary. The neighbourhood LAN model should allow current multicast mechanisms to be employed. Such traffic is normally at an acceptable level in the CPN but the exact dynamics of how this traffic may affect the many users of a neighbourhood LAN need further investigation. It may be that users are required to pay a premium to the ISP or CATV network operator to have access to the multicast address space. So, resource control mechanisms and packet filtering (e.g. based on firewalls) will be very important to cable operators offering such IP services.

In the ATM model, there are some fundamental differences in the group communication model which need to be addressed and these are summarised in [27]. A solution to these issues is under consideration [28] and it proposes the use of a multicast address resolution server (MARS) and (optionally) a multicast server (MS) to support the Internet model over ATM.

IPv4 has no native resource control or QoS mechanisms (F5). IPv6 specifications are more “resource aware”. For instance, the IPv6 packet header now contains a 24bit flowID field (the same size as an ATM VPI/VCI[6]) which could, for example, be used to distinguish packets that contain different types of media to implement a class-based QoS control mechanism. However, IP has no direct signalling facilities which can be used to explicitly ask for resources. Instead, the Internet model uses additional protocols to allow resource reservation and QoS control. When compared with ATM, the fundamental difference is that ATM uses the same protocol – Q.2931 – to set up a VC and to signal the QoS requirements of the VC, whilst the Internet model signals the QoS requirements separately from the set-up the information flow. A comparison summary of ST2+, RSVP and ATM is presented in RFC1821 [24]. Other signalling issues are discussed in [29].

The implementation of F5 in the DAVIC model would require further investigation. With a

[5] This is assuming that each STB has it’s own IP address – this does not have to be the case.

[6] The use of this field, however, remains a point of some discussion within the Internet community.

central server based approach, it should be easy to control the admission and servicing of resource/QoS requests. However it is currently unclear how these might be signalled to the subnetwork level, for instance, by the use of DSM-CC signalling.

Resource/QoS control for the neighbourhood LAN would require the use of additional packet classification and scheduling facilities within the network. Such mechanisms could be combined with traffic policing or traffic shaping functions that represent the provider's traffic/QoS policies. However, it may be very computationally expensive to provide this on a per subscriber basis in the neighbourhood LAN environment.

V. Summary

The great proportion of residential users must connect to their ISP by use of modems over analogue PSTN lines. Cable connections are currently coaxial cable based links that have a much higher raw bandwidth with the potential to offer high data rate services to a subscriber, particularly in the local area. From the discussion above, it can be seen that the items F1 to F5 give rise to the following connectivity requirements:

- **a suitable network integration interface for the CPN:** the interface at the CPN should be easy to use by the subscriber and should not (ideally) significantly change (from a users point of view) the current model of IP internetworking which is well understood and widely used. This requires the maturity of suitable standards, mainly from the IETF, ATM Forum and IEEE.
- **a familiar service offering:** users may have differing requirements of the kind of service that they require. A residential user will probably not require the same level of connectivity as a SME or corporate user. The interface at the CPN should allow all subscribers to make service requests that change the service they use in some controllable manner. IP users would like to use mechanisms like RSVP to adjust the connectivity that they receive.
- **addressing and routing should follow current and evolving Internet models:** the way that the provider network is configured should not affect the way in which subscribers choose to set-up their LIS in the CPN. The provider network should provide a transparent service, with appropriate address resolution and routing facilities. This may require additional resources in the provider network which perform address mappings or information transforms as necessary.
- **secure communication:** facilities to allow protection of data in transit and site-based access control (firewalls) are required. Such services may not be required by all users.
- **full Internet services:** the Internet user should be able to make use of services such as multicast communication and resource reservation through the commonly used interfaces and mechanisms. The network provider may have to introduce measures to police and control the use of such facilities within the network to ensure fairness and availability of network capacity to all users.
- **high speed, low latency communication:** Internet users expect (in some cases need) a high speed, low latency service to allow real-time interactive applications to function.

The use of ATM in the provider network, coupled with suitable IP-based protocols, provides the most attractive route for the provision of IP network connectivity if F1, F2, F4 and F5 are to be realised (F3 can be realised at the IP level using IP-based security services). The DAVIC framework allows for the use of ATM in the provider network, although the access to ATM interfaces for an IP user needs to be clarified. There is much work within the Internet community in providing IP over ATM, as well as work concerning the realisation of F1 – F5. It seems sensible to use standards from the Internet community in implementing Internet services, wherever possible. If a LAN-like interface is required for the CPN, then work from IEEE 802.14 and the ATM Forum should be considered.

The neighbourhood LAN model is a simple model that is relatively easy to implement. It can provide the CPN with an interface that is familiar to many IP users, but may not be able to provide F3 and F4 without significant effort or additional resources employed within the provider network.

The DAVIC model, using MPEG-2, has the advantage that it conforms to the DAVIC framework. However the value of this feature for providing Internet services is not clear. The DAVIC model should allow the provision of F1 and F2 but the DAVIC architecture is still maturing and the role of IP is unclear. The provision of F4 and F5 within this model has not yet been investigated but could be technically challenging at best, impractical at worst.

At the time of writing, the IPCDN model is being developed and, as it is an IETF activity, it will be very significant to the IP community.

The technology exists to provide a broadband local Internet connectivity, for instance in the form of a **neighbourhood Intranet**.

VI. References

[1] P. Heywood, "ISDN: A Long Time Coming … But Still No Quick Fix", *Data Communication International*, vol24(15) pp48A-48G, Nov 1995. McGraw-Hill.

[2] J. Barnes, J. Crowcroft, D. Farinacci, D. Ginsberg, D. Newson, D. Pratt, "Internet 2000: High performance support for real-time services", unpublished 1996.

[3] R. Braden, J. Postel, "Requirements for Internet Gateways", RFC1009, Jun 1987.

[4] The ATM Forum Technical Committee, "LAN Emulation Over ATM, Version 1.0", af-lane-0021.000, ATM Forum, Jan 1995.

[5] R. Cole, D. Shur, C. Villamizar, "IP over ATM: A Framework Document", RFC 1932, Apr 1996.

[6] M. Laubach, "Classical IP and ARP over ATM", RFC1577, Jan 1994.

[7] R. Braden, L. Zhang, D. Estrin, S. Herzog, S. Jamin, "Resource ReSerVation Protocol (RSVP)", work in progress, IETF, 1996.

[8] L. Degrossi, L. Berger, "Internet Stream Protocol version 2 (ST-2) Protocol Specification – Version ST2+", RFC1819, Aug 1995.

[9] M. Laubach, "Using ATM over Hybrid Fiber-Coax Networks", presented at the *Society of Photo-Optical Instrumentation Engineers (SPIE) Eastern Conference*, Philadelphia, PA, USA, 23 Oct 1995.

[10] DAVIC Guidelines #37, "Guidelines for Internet Access", work in progress, DAVIC Geneva, Switzerland (12b37r10.doc), 7 Mar 1996.

[11] ISO/IEC DIS 13818-6, "Information technology – Generic coding of moving pictures and associated audio information – Part 6: Extension for Digital Storage Media Command and Control", Nov 1995.

[12] M. Ahmed, G. Roeck, "IP Over Cable Data Network Service", work in progress, 1 Oct 1996

[13] J. Case, C. Davin, M. Fedor, "Simple Network Management Protocol – SNMP", RFC 1098, 1 Apr 1989.

[14] J. Case, K. McCloghrie, M. Rose, S. Waldbusser, "Protocol Operations for Version 2 of the Simple Network Management Protocol (SNMPv2)", RFC 1905, 22 Jan 1996

[15] M. Ahmed, K. Tesink, "Definitions of Managed Objects for ATM Management Version 8.0 using SMIv2", RFC 1695, 25 Aug 1994.

[16] R. Kosak (ed), "M4 SNMP MIB Network element", work in progress, ATM Forum, 1996.

[17] DAVIC 1.0 Specification Part7, "High and Mid Layer Protocols", DAVIC, Geneva, Switzerland, 1995.",

[18] D. Robinson, "DAVIC Server MIB for inclusion in part 7 Annex B", output from DAVIC, Seoul 4-8 Mar 1996, DAVIC (11b12r10.doc), 7 Mar 1996

[19] J. Postel, "Internet Protocol DARPA Internet Program Protocol Specification", RFC 791, Sep 1991.

[20] S. Deering, R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 1883, Dec 1995.

[21] Y. Rekhter, T. Li, "An Architecture for IPv6 Address Allocation", RFC1887, Dec 1995.

[22] H. Schulzrinne, S. Casner, R. Frederick, V. Jacobson, "RTP: A Transport Protocol for Real-Time Applications", RFC1889, Jan 1996.

[23] R. Braden, D. Clark, S. Shenker, "Integrated Services in the Internet Architecture: an Overview", RFC1633, Jun 1994.

[24] M. Borden, E. Crawley, B. Davie, S. Batsell, "Integration of Real-Time Services in an IP-ATM Network Architecture", RFC1821, Aug 1995.

[25] R. Atkinson, "Security Architecture for the Internet Protocol", Internet RFC 1825, 9 Aug, 1995

[26] M. Leech, M. Ganis, Y. Lee, R. Kuris, D. Koblas, L. Jones, "SOCKS Protocol Version 5", RFC1928, 3 Apr 1996.

[27] D. Ginsburg, J. Barnes, D. Newson, D. Pratt, "IP multicast distribution of real-time MPEG over ATM", *JENC7 – 7th Joint European Networking Conference,* 13-16 May 1996, Budapest.

[28] G. Armitage, "Support for Multicast over UNI 3.1 based ATM Networks", work in progress, IETF, 1995.

[29] M. Perez, F. Liaw, A. Mankin, E. Hoffman, D. Grossman, A. Malis, "ATM Signalling Support for IP over ATM", RFC1755, February 1995.

Author Information

Saleem N. Bhatti holds a B.Eng.(Hons) in Electronic Engineering (1990) and a M.Sc. in Data Communication Networks and Distributed Systems (1991), both from University College London. Since October 1991 he as been a member of the research staff at UCL, involved in post graduate teaching and working on various computer communication and networking related projects, covering topics such as network and systems management, network security, ISDN, ATM, QoS, and residential broadband services. Saleem is a member of the IEEE and the ACM.
http://www.cs.ucl.ac.uk/staff/S.Bhatti/

Graham Knight graduated in Mathematics from the University of Southampton in 1969. He taught Maths and Computer Science in schools and colleges until 1979 when he obtained an MSc in Computer Science from University College London. Since then, he has worked at the college, initially as a Research Assistant, now as a Lecturer. His research work has been in the field of computer communications. He has led UCL teams in ESPRIT projects concerning OSI, Network Management and ISDN, and is now involved in ACTS. Currently his research interests include routing in hybrid SDH/ATM networks, CATV-based Internet access and the use of Java and WWW for network and systems management.
http://www.cs.ucl.ac.uk/staff/G.Knight/