Travertine Systems, Inc.

Product and Technology Overview
V0.6

8/14/97
Pete Foley
Updated 9/24/97

Table of Contents

List of Figures

List of Tables

Introduction

This document contains an overview of the Travertine Home Area Network (HAN). This overview is a working document that is intended to provide the reader with background as to the product goals and technical motivation behind the HAN, as well as to a smattering of the implementation details. This document does not explore the business development/OEM endorsement issues that must be addressed in order to realize the stated product goals.

All of the novel concepts presented below are covered in patents that Travertine has either filed, or is in the process of filing.

Product Motivation

With the explosive growth in the number of second PCs being purchased for home use (in excess of 30% of all household PC's now purchased are "additional" PCs), Travertine believes there is significant pent-up demand to network these PC's for purposes of resource sharing, multiplayer game playing, WAN access, and digital video distribution. This demand will only continue to grow.

Kim Maxwell, a founder of Amati and chairman of the ADSL Forum, believes there is an $8B market, akin to the home burglary alarm installation market, for pulling new Unshielded Twisted Pair (UTP) cable in existing homes to support home networking. Travertine believes otherwise. Travertine believes that the average consumer is not a very good systems integrator and is averse to the expense and inconvenience of having a third party enter the home to pull new cable. Furthermore, after many years and many dollars spent by the Telcos, those same consumers are very comfortable with the connectivity paradigm of the ubiquitous (and simple!) RJ-11 phone jack.

Organizations such as the ADSL Forum and the VESA Home Networks committee are advocating new premises hub and spoke installations using CAT-5 UTP wiring, with the local loop terminating at a premises hub/gateway. Besides the cost and inconvenience of new cable and an active hub/gateway installation, this type of architecture creates hub placement, hub power, environmental (thermal, etc.), lifeline POTS service, and other issues.

Another architectural approach is to unify the local loop and premises wiring under a common physical layer and protocol, which likely requires the installation of new equipment at Telco central offices. This approach may or may not subsume POTS. Unifying the subscriber loop and the premises wiring presents unique challenges in that the subscriber loop and premises wiring domains place quite different design and regulatory constraints on the systems designer. For example, the 12kft subscriber loop served by the Carrier Service Area (CSA) does not support signals much above 1.1Mhz due to strong attenuation, and the noise and crosstalk environment of the subscriber loop is quite different from that of premises wiring.

While this unifying approach can offer benefits in terms of simplifying the integration issues between the premises LAN and the subscriber loop WAN, the Telcos have historically proven to be quite conservative, slow to adopt new technology, and averse to the installation of new loop plant. From a marketing and product development standpoint, this approach requires bridging the disparate worlds of the Telcos and the PC/consumer electronics OEMs. Attempting to sell line cards to the Telcos may not be a good business for a startup because of the long runway, and appropriate partnerships would be a key factor in the success of such a business plan. Just dealing with the additional power requirements of DMT (vs CAP) ADSL modems has proven a barrier to installation of DMT ADSL modems in the central office. Furthermore, it is estimated that by the year 2000, only 10% of the homes in the US will have WAN access. There are also many players in the subscriber loop technology marketplace, with ADSL, ISDN, HDSL, FTTC, FTTN, and FTTH systems and providers all vying for prominence. There are also many other possible non-subscriber loop broadband suppliers that will be providing internet access and/or other services traditionally provided through the local loop. These include cable, terrestrial microwave (MMDS, etc.), and satellite sources.

On the other hand, if inexpensive CMOS radios in the 2.4 Ghz band should become available in the near future, then the ability to offer "seamless" connection to the subscriber loop could be a key factor in the overall value proposition.

Travertine has chosen to focus strictly on the home, and more particularly on the use of the existing in-home UTP phone wiring as the application space for its LAN technology and products, while maintaining interoperability with as wide a variety of broadband sources (to the home) as possible, but most particularly with traditional POTS, and with ADSL. This approach possesses the virtues of focus, of minimizing regulatory issues, of having a ready market, of having an existing infrastructure, aligning well with the goals of Intel and Microsoft, and of providing some independence from the Telcos and the broadband provider/standards competition. This product focus is PC/consumer electronics centric, rather than Telco centric.

Product Description

The Travertine HAN enables premises broadband networking using the existing installed POTS Unshielded Twisted Pair (UTP) wiring. Devices are attached to the HAN in the same manner that phones are attached to the PSTN, by simply plugging into an existing RJ-11 wall jack.

By using frequency division multiplexing to take advantage of the underutilized spectrum from approximately 1 to 20Mhz, the HAN is compatible/interoperable with existing POTS and ADSL standards and requires no changes to the loop plant infrastructure and no new assumptions regarding subscriber loop signaling protocols.

The HAN anticipates and is capable of delivery of broadband digital data throughout the home from such varied sources as subscriber loop (xDSL), satellite (DSS/DVB), terrestrial microwave (MMDS), digital cable/CATV ("All TV"), digital or high definition television (DTV/HDTV/ATV), and Digital Video Disk (DVD) drive.

The HAN facilitates "convergence anywhere" in the home.

Figure 1 shows broadband networking within the home as it typically exists today – using Ethernet coaxial cable and/or 10/100baseT UTP wiring in a star configuration concentrated at a hub. Note that in the case of 10/100baseT, a new standalone active device is typically required – the hub – and that new UTP cabling must be pulled through the premises to network the computers in point-to-point (star) fashion.

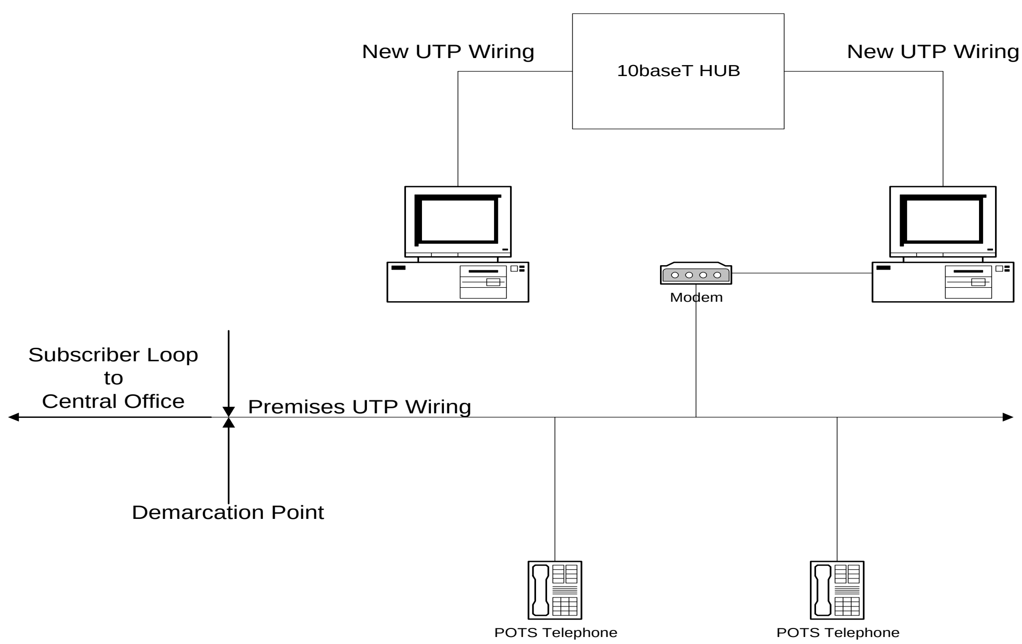

Figure 1: Existing Premises Networking

New proposals for home networking require installing new CAT-5 UTP cable, typically in a point-to-point star fashion centered about a home gateway. Such a topology allows baseband modulation schemes to be used, and these proposals target 50-100Mb/s rates using relatively simple equalization.

However, new cable must be pulled and an active hub installed. Such a system has been put before the VESA Home Networking committee, for example.

Figure 2 shows the same functionality achieved using the existing phone lines employing the Travertine HAN technology. The familiar RJ-11 plug-n-play paradigm is employed for both computer networking and for POTS access. The user need not be concerned about network topology during installation.

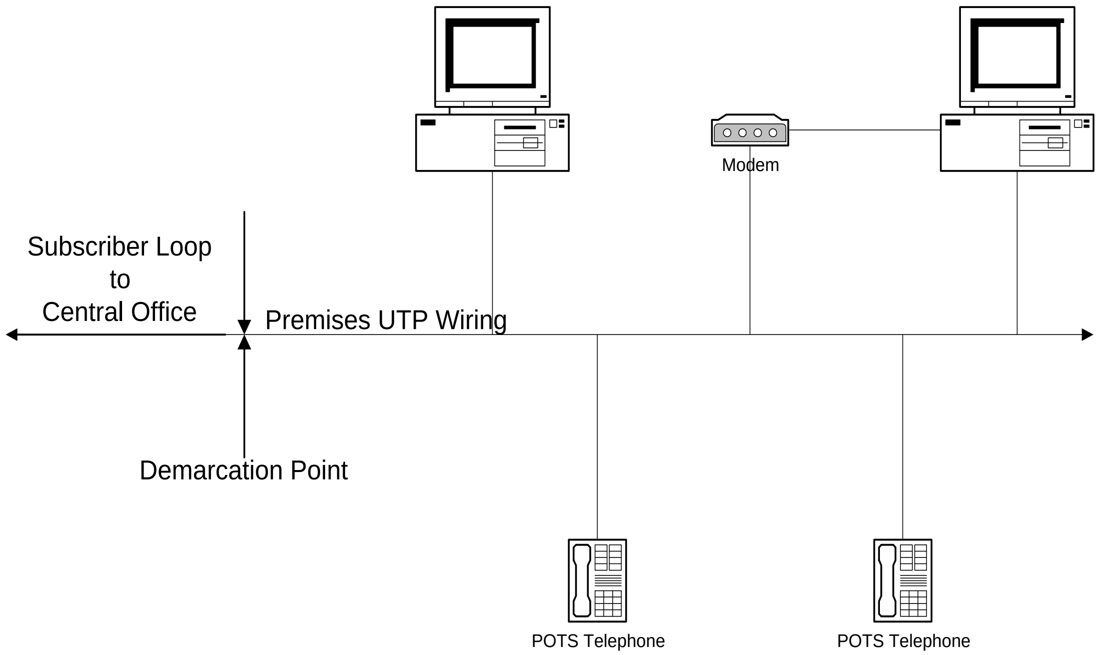

Figure 2: Premises Networking Using Travertine HAN

The Travertine HAN can alternately be described, from a marketing standpoint, by the term **ResiNet**™, which can be an acronym for *Residential Network*, or for *Residential Ethernet*. Travertine originally used ResiNet as its company name, but because of potential conflicts with similar names changed to Travertine Systems. Travertine owns the WWW domain ResiNet.com, and is still pursuing the possibility of obtaining trademark status on the term ResiNet.

Travertine is investigating obtaining trademark status for the term **100baseH** (100baseHAN) to apply to the HAN as well.

Product Goals

Long Term Business Goals

A long term business goal of Travertine *is to integrate the Travertine HAN technology into every standard POTS modem chipset sold*.

As the HAN becomes established in the marketplace, Travertine will migrate the business to an IP/royalty model for the core technology, while focusing on designing products utilizing the HAN.

Product Features and Benefits

In order to achieve the desired long term goals, the Travertine HAN technology must become ubiquitous, and this ubiquity is greatly facilitated via the following product requirements:

a) support of arbitrary premises wiring topologies,
b) support for variable wiring quality,
c) support of legacy POTS devices (modems, phones, etc),
d) plug-n-play paradigm (trivial insertion/removal of new network devices),
e) robust operation in the face of a variety of ingress sources,
f) graceful degradation when faced with exceptionally hostile channel characteristics,
g) compliance with existing regulatory requirements

If these requirements can be achieved, then Travertine can enable WAN/LAN integration to occur in the home with the following benefits:

a) using the existing well accepted POTS plug-n-play paradigm,
b) while maintaining compatibility with existing POTS services, including xDSL **and lifeline service** (supporting DC current flow from the Central Office (CO)),
c) using the existing installed POTS UTP wiring infrastructure,
d) without requiring the use of active gateway/bridge devices,
e) while providing broadband LAN performance.

Robustness

It is critical to the acceptance of the use of premises UTP wiring as a network infrastructure that the technology employed be perceived as extremely robust (trustworthy) by the marketplace. A bad reputation is slow to repair and the home market is a fickle one. The reputation of RF networking, for example, as having interference and security problems has hindered its acceptance both by the marketplace and by the financial/venture community.

A significant risk to the establishment of a HAN market using existing UTP wiring would be the premature introduction of products that that are technically inadequate. Travertine believes that products that do not focus exclusively on the home networking problem, but instead try to solve a broader set of requirements, in particular run this risk. Travertine’s definition of technical inadequacy includes systems that do not possess sufficient margin to service the overwhelming percentage of homes despite the hostile environment presented by premises wiring.

A field “failure” rate of even a few percent could have disasterous consequences on product reputation, support costs/infrastructure.

The subscriber loop, which is a far more controlled environment than the premises, has been extensively studied and is well characterized and modeled (particularly by Bellcore). Despite this, adoption of both ISDN and xDSL technology has been slow, principally because the RBOCs have discovered that an unacceptably large fraction of the installations require either the installation of new loop plant or significant characterization time by trained personnel.

Unlike the subscriber loop, very little data is available on the characteristics of premises wiring, and little research has been published. Anecdotal reports are generally that the premises UTP wiring environment is a very hostile one. There have been no large (public) surveys made from which to extrapolate the necessary requirements for coverage of the majority of homes. And unlike the subscriber loop, the premises wiring network is not only more hostile because of the much greater topological variability, splicing, and wiring quality, but the network can also change dynamically. Network characteristics can change by actions as simple as taking a phone off hook, or as significant as plugging in a new bridged tap (stub) into the network.

Because of these reasons, and because Travertine's premises wiring surveys/analysis to date have been necessarily modest, Travertine has selected the most conservative/robust modulation scheme (provides the greatest Signal to Noise Ratio (SNR) headroom) we believe is available in off-the-shelf form – that still meets the data rate goals. That modulation scheme is QPSK (or QAM-4). When performing a link budget analysis for a reasonably well characterized environment like the subscriber loop, the designer can specify a modest SNR margin to hedge against contingencies – however because of the much greater variability of premises wiring, the margin specified must be commensurately greater. The greater the margin/robustness, the greater will be the coverage of existing homes, and the fewer the product returns/support issues. Opportunities for cost reductions by relaxation of requirements for BALUNs and other AFE interface logic are also facilitated by sufficient SNR margin.

Data Rate

The primary performance driver of the HAN is digital video, specifically DSS/DVB/DVD MPEG-II data streams. Current MPEG-II data streams typically consume 5 Mb/s on average, with sustained peaks in the 9 Mb/s range for DVD. DirecTV licensees are required to support the full MPEG-II main-level main-profile data rate of 15 Mb/s, although this rate is rarely seen actual use. Advanced Television (ATV) standards calls for a data rate of up to 19.2 Mb/s.

Travertine wishes to support either two simultaneous MPEG-II data streams, or a single ATV data stream, along with the paging activity associated with remote Windows 95 clients. The latter will provide for classic client-server services and NetPC type products. These requirements have driven the Travertine HAN sustained data rate goal of **20 Mb/s**. *(Anecdotal reports from an RF LAN vendor indicate that the consumer electronics companies want LAN's in the 20-25Mbit/s range).*

Product Applications

Applications for the HAN technology include:

a) Resource sharing – such as printers, scanners, fax machines, hard disk drives, DSS receivers/set-top boxes, DVD drives, modems, broadband WAN access devices such as xDSL modems, etc.
b) Web browsing using stations remote from the Internet/WAN access device
c) Multi-player gaming, either in the home or over the Internet
d) Distribution of digital video and/or audio throughout the home
e) Support for client/server applications and products, such as NetPCs

Network Latency

Minimizing network latency is a goal of the Travertine HAN in order to support paging over the network and to support latency sensitive applications such as multiplayer Internet gaming.

Internet gaming can be particularly latency sensitive for so called "twitch games". Each game action typically requires two round trips between players – or four modem transits. Because of modem transit latencies, Internet game play often has to be regionalized when signal flight times and other network delays are taken into account, and any latency the HAN additionally adds will further shrink the coverage radius. Since multiplayer gaming is a large and growing market, and a principal draw for teenagers wishing second PC's, low latency is a HAN priority.

Cost Goals

To minimize the barrier to market entry, Travertine will develop a solution that has a connect cost goal of ~$25 per node in volume. This requires the development of an integrated single chip HAN modem ASIC with a high volume cost goal of < $12. This should be quite feasible in .25 micron technology.

The HAN connect cost will be further reduced – to be essentially "for free" – as the HAN technology is integrated with the standard POTS modem.

Initial Products

The HAN technology will be initially brought to the market through the simultaneous introduction of two vehicles, the HAN Network Interface Card (NIC) (the “HANic”), and the HAN USB dongle (the “HANster”). These products will have a sustained data rate in the 6 Mb/s range (the NIC card may support higher data rates), sufficient to support a single MPEG-II data stream.

Travertine intends to develop reference designs for these products to facilitate their availability to PC OEM’s through the OEM channel. If necessary, Travertine may manufacture NIC cards and dongles in order to “jumpstart” market entry, along the lines of the WebTV model.

Figure 3 shows a block diagram of the first HANic using off-the-shelf CATV modem technology.

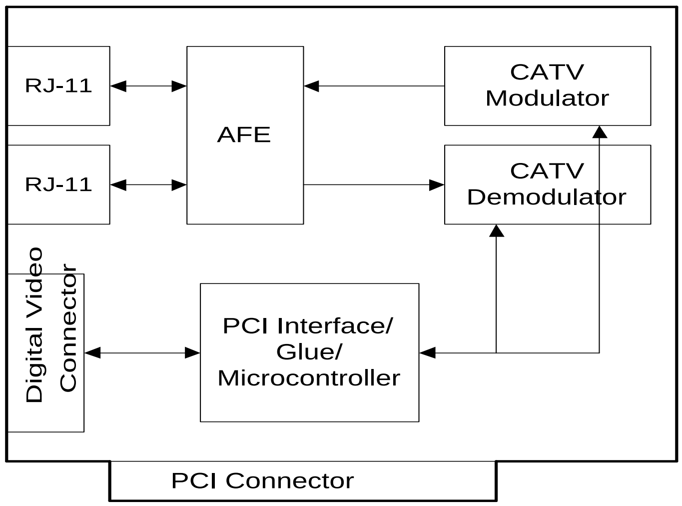

Figure 3: Initial HAN NIC implementation

This PCI add-in card uses off-the-shelf CATV modulator and demodulator chips (single chip integrated CATV subscriber modems are becoming available, and it may be posible to use them instead), along with a PCI interface logic and the glue logic required to interface to the CATV chips. A local microcontroller may be required to support aspects of the HAN protocol (such as equalizer coefficient storage) if host latency proves to be too great. Although not strictly required, two RJ-11 jacks are shown to allow connection between separate UTP lines in the premises (allowing the two UTP lines together to support a single network). The digital video connector allows digital video to be gated onto the HAN without first passing through the host CPU or host PCI bus. The initial target for the Digital Video Connector is the Wideband Data Port (DB-25 connector) found on satellite set-top boxes (this data port is moving to FireWire in future versions).

Figure 4 shows a block diagram of the HANic using the integrated HAN modem chip.

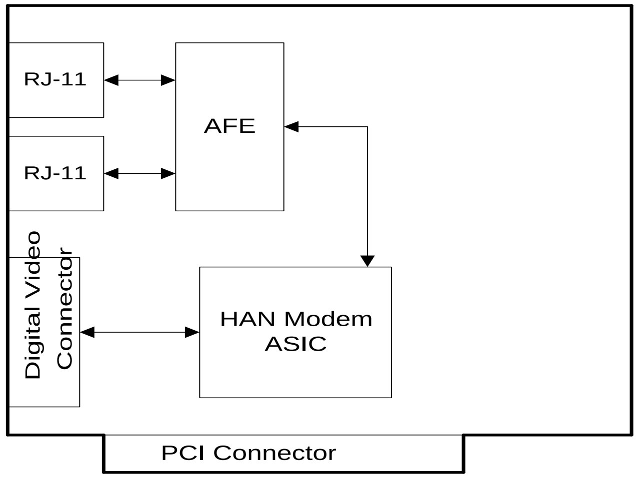

Figure 4: HAN NIC using integrated HAN modem ASIC (2nd generation product)

Besides significantly lowering product cost and footprint, the integrated HAN modem ASIC will allow the HANic to operate at the full data rate goal of 20 Mb/s.

Figure 5 shows a block diagram of the initial version of the HANster USB dongle.

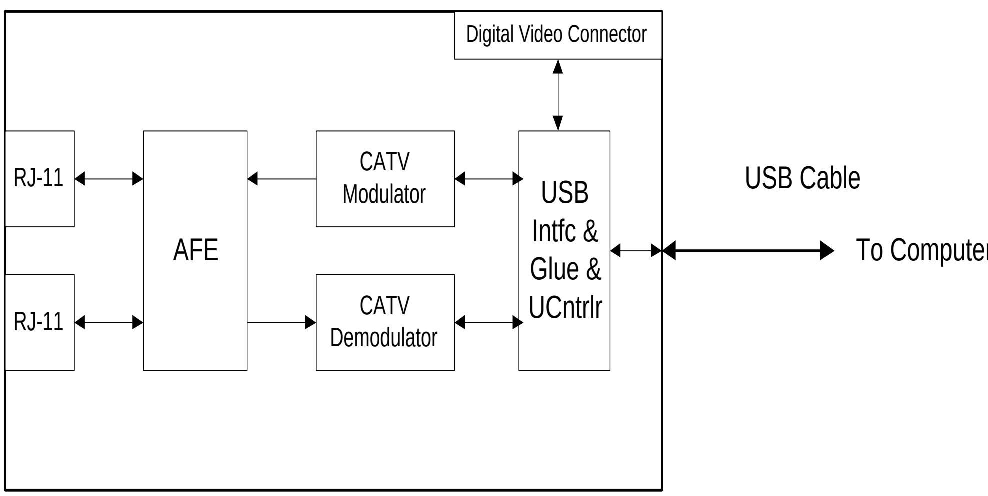

Figure 5: Initial HANster product implementation

As with the NIC card, a microcontroller may be incorporated if host control of the HAN network introduces excessive latency. Incorporation of a local microcontroller is more likely in the case of the HANster because of the additional communications latency (to/from the host CPU) imposed by the USB. The performance of the USB HANster will be limited by the sustained USB throughput for streaming data of ~ 6 Mb/s.

Figure 6 shows the second generation peripheral "dongle" HANster using the integrated HAN modem. In this implementation, the throughput limitations are eliminated both through the use of the integrated HAN modem ASIC, and through the migration of the host interface to the 100Mb/s FireWire (IEEE-1394) instead of USB (Set-top box wideband data ports are expected to migrate to FireWire as well).

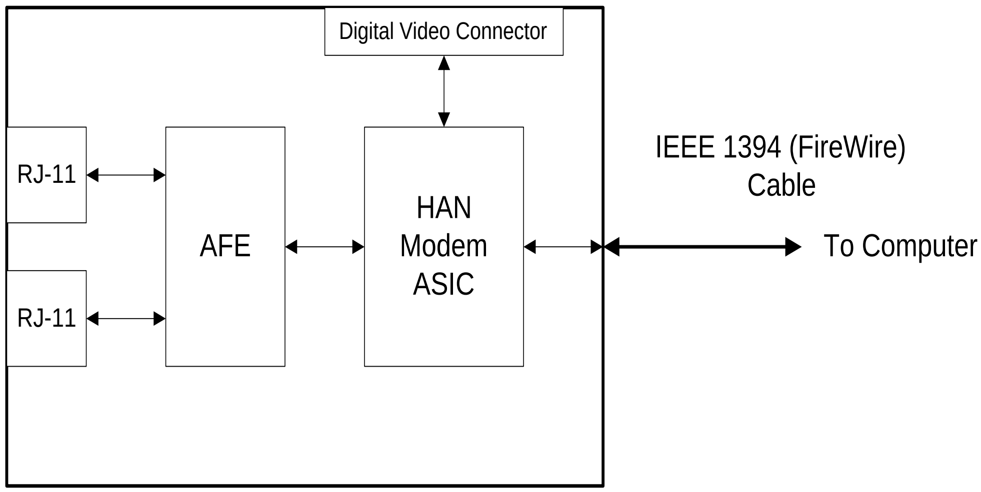

Figure 6: HANster implementation using HAN Modem ASIC (2nd generation product)

Follow-on Products

A variety of follow-on HAN based products are possible:

Motherboard Licensing/SecondPC™

Once the NIC and dongle products are in the market, a subsequent revenue stream is expected from the licensing of HAN technology for use on PC motherboards.

Low cost network computer/client computer products which leverage expensive resources in a standard client/server model become attractive as the HAN technology is integrated onto the motherboard. Travertine has applied for trademark status for **SecondPC** to describe these low cost, potentially diskless, client PCs (NetPCs). Travertine also owns the WWW domain SecondPC.com. It should be noted that unlike NetPCs that are Java based, these SecondPCs could be full fledged Windows 95 machines.

Telepliance™

As the HAN's market adoption grows, a dizzying variety of network appliance products become viable, including low cost web browsers, answering machines, remote video cameras, and remote audio speakers. Travertine has applied for trademark status for **Telepliance** to describe such products. Telepliance could be used as a generic term, or could be used for specific Travertine products, whichever proves most advantageous from a marketing standpoint.

Teleplayer™

Another interesting product possibility is a "universal" portable MPEG-II player device – the **Teleplayer**. This digital video player concept would be universal both in its ability to be located anywhere in the home, and in its ability to decode a variety of digital video formats. This product would allow the playback of MPEG-II video from any source, satellite, DVD ROM, digital cable, MMDS, ATV, etc. located anywhere in the home. The Teleplayer could be plugged into an RJ-11 jack wherever it is desired to playback the video. This product would consist principally of the integrated HAN modem ASIC, and an integrated DVD/MPEG-II decoder and an NTSC/PAL encoder. This product could be a unique differentiator for the HAN because of the high bandwidth requirements.

Virtual Gateway™/WAN access integration

WAN access integration presents possibilities for software only as well as integrated hardware/software products.

In the case of ADSL, a software product which performs typical bridge functions such as protocol translation, buffering, forwarding, etc. is possible to facilitate WAN access to/from the HAN. Because this WAN/LAN bridging can be realized in software on a host PC without the need to break the subscriber loop as it arrives at the premises – in other words the subscriber loop need not terminate at an active device in the premises – the term **Virtual Gateway** is used to describe this technology. Additional detail on the Virtual Gateway product implementation are described later in this document.

A possible hardware product would be the integration of ADSL and HAN modems onto a single PCI NIC card. The only connector(s) needed on the card edge for this product is an RJ-11 jack. The Virtual Gateway/Bridging software would be bundled with this add-in card.

HAN Technology Description

Subscriber Loop vs Premises Wiring

Given the close relationship of the HAN and the subscriber loop, some of the salient characteristics of both are presented:

Subscriber Loop

Subscriber loops of 18Kft encompasses 80% of US subscriber loops. Due to strong signal attenuation (up to 90dB at this cable length), realistic downstream rates of only 1.5Mbit are possible.

The Carrier Serving Area (CSA) is defined as subscriber loops of 12Kft or less, which covers 50% of US market. For this cable length, 6Mb/s is a feasible ADSL downstream payload.

Average length of residential loops is 7865 ft according to the 1983 Bell System Loop Survey.

15-20% of US residential lines have loading coils - these cannot be served by xDSL

The average subscriber line has 22 splices and 3.7 bridged taps with an average length of 1186 ft. A bridged tap is an unterminated spur that acts like a delay line and puts a notch (spectral null) in a line's frequency/attenuation characteristic at the frequency associated with a bridged tap's characteristic wavelength.

xDSL usable bandwidth is usually limited by Near End Crosstalk (NEXT) originating from other xDSL or ISDN transmitters that couple to the subscriber loop in the cable bundle outside the premises.

NEXT and FEXT usable bandwidths are defined by cable and crosstalk losses, while the white noise bandwidth is expandable by raising the transmit power density.

Inter-pair crosstalk increases as freqency raised to the 3/2 power, and increases as the number of adjacent disturbers (in the cable) raised to the .6 power.

Home Wiring

Typically, low-quality nontwisted biquad is used for home wiring. Two balanced wire-pairs exist in the biquad which are usually color coded.

Virtually any wiring topology can exist in the home, including daisy chain, tree, and star topologies. The extent to which loop topologies are present is TBD.

Home wiring lengths vary widely, with overall lengths probably longer than one might expect, particularly for larger multi-story dwellings. Even small (1600 sq. ft.) single story homes can have premises wiring length from the Telephone Network Interface (TNI) of about 450'. The longest peer-to-peer connection length will be some fraction (expected to vary from 0.5 to 0.9) of the overall premises wiring length. However, since many homes contain two independent phone lines, *Travertine is proposing a technique for "shorting together" multiple lines in the home to allow these lines to serve as one physical network in the spectral domain of the HAN* (this is explained in more detail later in this document). Depending on where the two lines are "shorted" there is a likelihood this practice will increase the maximum peer-to-peer connection length. So for the sake of analysis, and until much more empirical/survey data is gathered, Travertine's current conservative channel length target is 1000 ft. Again

Travertine emphasize the importance of designing a conservative system that will service the majority of homes.

At a length of 1000 ft, the typical premises biquad wiring will exhibit a rapid increase in attenuation with frequency, such that at 18Mhz, approx. 50dB of attenuation can be seen if there are no taps on the line and no Customer Premises Equipment (CPE) is plugged in (ie a clean line is examined). The effect of short bridged taps (taps under the length of 8') is to increase the attenuation slope. Attachment of a single 3' bridged tap will increase the above attenuation 10 dB (to ~60dB) at 18Mhz.

Bridged taps longer than 8' start to cause visible spectral nulls or notches in the cable transmissivity at frequencies below 20Mhz. An 11' silver satin extension, for example, causes a 20dB notch at 16Mhz. These notches can be compensated for by proper equalization.

POTS/xDSL Compatibility

From a marketing standpoint, not interfering with POTS and secondary digital services is as important as network robustness. The user must not perceive a reduction in quality of service associated with these functions when using their installed UTP for the HAN.

As was previously mentioned, POTS and ADSL compatibility is maintained by utilizing spectrum above 1.1Mhz. Figure 7 shows the spectral allocation for a typical HAN implementation.

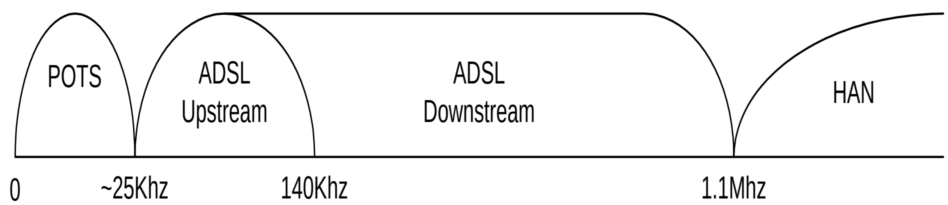

Figure 7: Typical HAN Spectral Allocation

In an actual implementation, the start of the HAN passband would be located sufficiently far above the upper edge of the ADSL downstream band so as to facilitate the use of low cost passive filters to isolate the two bands. Typical proposed starting frequencies for the HAN passband are in the 3-4 Mhz range. Insertion of a sufficiently large transition zone between the two bands allows the use of filters with reasonable skirts (lowering cost) and allows the movement of filter cutoff frequencies into the transition zone to keep group delay peaks and other undesirable filter characteristics out of the passbands.

A single passband implementation of the HAN is intended to support half-duplex connections (like Ethernet). The allocation of an additional band(s) allows full duplex connections between peers and/or allows the creation of another completely independent network of peers. Movement between these "networks" can be supported if the network client HAN modem implementations are frequency agile.

HAN Spectral Allocation Considerations

Figure 8 shows the spectral allocation for a multi-band HAN implementation

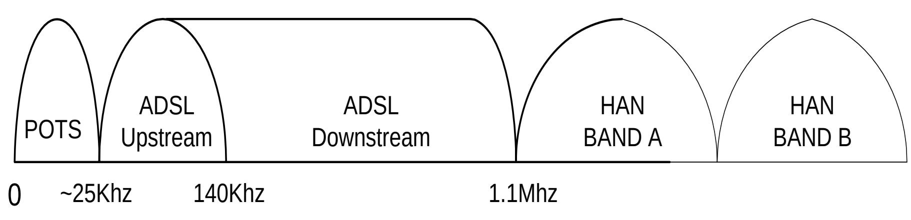

Figure 8: Multi-band HAN implementation

Practical limits on the upper frequency expected in a HAN system limit the number of bands possible if the per band spectrum allocation does not decrease. For frequencies much above 20Mhz, biquad UTP attenuation within the home (for the target 1000 ft length) and radiated emissions become the constraining factors. There is also the issue of the passive user installed filters designed to improve HAN performance. These filters, in order to pass DC current and to support signaling currents/voltages required by CFR Part 68, use wire-wound inductors that self-resonate at frequencies much above 20 Mhz.

For these reasons the target HAN spectrum allocation is roughly 16 Mhz, from ~4 Mhz to ~20 Mhz. QPSK or QAM-4 signaling then provides 32 Mb/s of raw payload, or ~28 Mb/s of payload after Reed-Solomon overhead (roughly 10-12%) is accounted for. This 28 Mb/s allows an additional 30% for protocol and other overhead while still meeting the ~20 Mb/s throughput goal. A 30% allocation for protocol and access overhead is aggressive. Given the known issues with CSMA/CD protocols such as Ethernet (probabilistic MAC protocols lead to uncertain media access delays, low throughput under heavy loading, etc), and the difficulties associated with practical carrier sense and collision detection in a passband modulation scheme, Travertine will likely adopt a non CSMA/CD access protocol. This will be explored later in this document. Note that the use of convolutional encoding as an additional error correction technique will further reduce the net data rate. Additional spectrum could be allocated, or potentially a higher constellation QAM (such as QAM-16) could be adopted if the net data rate needs to be increased and sufficient channel SNR exists. The principal method of increasing the channel SNR in this highly multipath dominated environment is to increase the temporal (number of baud times) coverage of the equalizer.

PSTN Regulatory Issues (CFR 47, Part 68)

Characteristic Impedance/Termination Impedance

The characteristic impedance of the cable is largely irrelevant until cable lengths and bit times approach the region where the behaviour of the EM wave on the cable is important. In a typical cable, the wavelength of the audio band is considerably longer than the subscriber loop. Hence the choice of line termination impedance for low frequency POTS is set by other factors such as noise and drive current availability; 600ohms has always been the traditional value.

Above 1Mhz the characteristic impedance of UTP is well behaved and is almost purely resistive with a value of about 100 ohms, and 100 ohms is the value Travertine uses for matching impedances for the HAN.

Allowed Transmit Power (on the subscriber loop)

A transmit power of 0dBm corresponds to a P-P voltage of .32V. Most of the modulator ASICs evaluated transmit with 1Vp-p max, which would allow a transmit power of 10 dBm (@ 100 ohm) without additional amplification. At the target baud rates, this is about the max power that is expected to still meet FCC class B.

Part 68 signal power limitations are considerably more stringent over the lower 250Khz of the ADSL range (below 270Khz) than they are for the HAN spectral range; with requirements that metallic voltage not exceed –55dBV (1.8 mV) in the 86 – 270 Khz range. This is a stringent requirement that hampers use of this spectral range for premises networking in that much higher power levels could be used if FCC class B were the only requirement. As a result, the use of a premises gateway or strong filter at the TNI to allow a higher transmit power (in the premises) in this spectral range would be required. Fortunately this issue does not confront the Travertine HAN because of the higher spectral domain chosen.

Part 68 requires that the metallic voltage not exceed –15dBV in the range from 270Khz to 6Mhz, and that the longitudinal voltage not exceed -30dBV in the same frequency range. The metallic restriction corresponds to a differential voltage of 0.178V, or a transmit power of -5dBm, which means that the TNI filter must supply 15dB of attenuation to guarantee Part 68 compliance if a 10dBm transmit power is used. Practically speaking, the UTP itself will provide 15dB of attenuation to the lower edge of the HAN passband after less than a thousand feet.

Radiated emissions result from conversion of differential/metallic signals into common-mode signals (currents) by electrical imbalances in the cable and at the transceiver ports. Such imbalances are a principal cause of susceptibility to induced noise as well. Simple BALUN termination may not give optimal results from an EMI standpoint. These issues are discussed in some detail in [1].

Although empirical verification is needed, Travertine believes that at frequencies below 20Mhz, a 10dBm power level will not violate FCC Class B limits with well filtered signals, and that in situations where a TNI filter is not installed, an option for transmitting at a 10dBm power level is to transmit at frequencies above where the Part 68 restrictions apply (above 6 Mhz). This is an advantage of placing the HAN band above the ADSL frequency range.

Insertion Loss

See the filter section.

Return Loss

The higher the return loss, the better the impedance match and the lower the reflected signal – which minimizes echo. Return loss specs play a principal role in the design of POTS splitters for ADSL modem AFEs, but don't play much of a role in the HAN filter designs because POTS is well down in the HAN filter passband.

Minimum return loss requirements looking into the subscriber loop from the CO is:

3.5dB	for 200Hz - 500Hz
7dB	for 500Hz - 2.5Khz
3.5dB	for 2.5Khz - 3.4Khz

Ringer Equivalence (REN)

A REN value of 1 implies 25Mohm DC resistance up to 100Vdc, and at least 150Kohm DC resistance for voltages from 100V to 200V. The REN of the HAN system is TBD.

Protection

Section 68.302 covers the metallic and longitudinal surge requirements. For the HAN filter designs, robust components are chosen in order to mitigate protection concerns. Usually a voltage rating of 300V is specified for anything always connected to the PSTN. A Metal Oxide Varistor (MOV) or similar device is used to shunt excessive metallic voltages when capacitors are employed as filter elements.

Part 15 Issues

There is little data available for radiated emissions (FCC class B) for existing premises wiring. There is some data developed for HDSL/VDSL systems that assume point to point installation of CAT-5 wiring [1]. For these systems radiated emissions spectra of UTP-5 were found to be surprisingly constant and independent of length and geometry. However there is a substantial rise in the emissions spectra above 40Mhz. Transmitted voltages of 100Mbps systems must be as low as .57Vp-p or less to satisfy FCC class B. 50Mb/s systems could transmit up to 3Vp-p without violating class B.

BALUN quality, particularly balance, along with appropriate termination structures, are important in insuring that no metallic imbalances are created that will result in increases in radiated emissions. Structural Return Loss (SRL) characteristics of UTP are an important determiner of radiated emissions. We have not seen any data relating typical home wiring practices to equivalent SRL factors and resultant radiated emissions. *The general issue of radiated emissions will be an area of early and intensive study by Travertine as it is regarded as one of the highest risk areas of the technology*.

Scrambling is generally used to disperse transmit signal energy throughout the band in order to prevent repeating patterns in the data stream from producing strong spectral lines. The HAN assumes scrambling is employed.

Modulation Scheme

QPSK/QAM-4 is currently the modulation scheme of choice because of its robustness and coding/encoding simplicity. With its simple encoding/decoding, QPSK lends itself to dedicated hardware solutions. Acquisition time is also lower for QPSK than for large constellation QAM systems, a consideration in a multi-peer network.

Because the spectral nulls due to reflections are frequency specific (tuned to the length of the bridged taps or stubs), spread spectrum techniques could be used to circumvent this problem, however

Travertine wishes to deploy a first generation product with off-the-shelf silicon, and nothing appropriate to the task was identified that could support the necessary data rate in the available spectrum.

In practical systems, QPSK requires about a 15dB SNR for reception at an acceptable bit error rate (BER), for an equivalent error rate, QAM-16 requires about a 25dB SNR, and each further increase in constellation size typically requires an additional 3dB in SNR. Travertine is targeting a BER of no lower than 1 in 10**-7 with 6 dB of margin (this is in line with BER targets of ADSL modems).

QPSK can utilize a 6-bit A/D for reception, whereas larger constallation QAM and DMT systems typically require 8 or 10 bit A/Ds. There are more "unknowns" or "degrees of freedom" in CAP/QAM and DMT systems than there are for QPSK that require greater discrimination, hence an extra 2 to 4 bits are required to provide the additional 12-24dB of dynamic range. These include:

a) Automatic Gain Control (AGC) uncertainty (need AGC for multilevel AM systems). QAM schemes, such as 16-QAM are slower to achieve AGC convergence, and so typically require longer preambles.
b) Amplitude and phase levels are much closer together. QAM systems, such as 16-QAM are more sensitive to InterSymbol Interference (ISI) caused by amplitude and delay imperfections in the channel. This in turn requires finer steps in the tap weights on the compensating equalizer network, and longer convergence times during equalizer training. Note also that received BER is dependent both on SNR and the amount of ISI, and ISI is largely independent of transmit power or other techniques used to improve SNR. Hence, for the same quality of equalization, QPSK will outperform QAM-16 from a BER standpoint for situations that are not SNR limited.
c) Crosstalk between In-Phase (I) and Quadrature (Q) components can result in constellation rotation, so are must be taking to insure the measurements are orthogonal w.r.t the I/Q space. The constellation points for 16 QAM are a factor of 3 closer together (amplitude and angle) than QPSK. This requires the system to have more SNR, less phase jitter, and less ISI.

Since all of the information is contained in only the phase of QPSK signals, as opposed to both the two dimensional coding (phase and amplitude) for QAM-16 and higher constellation QAM signals, the received signal can be passed through a hard limiter for recovery purposes, and this will add an additional 6dB of noise immunity for deterministic noise sources – such as CW emitters or power line AC pickup.

The use of 6-bit A/Ds greatly facilitates the flash sample rates required for the higher baud rate QPSK systems.

Current DSS/DVB satellite demodulators support QPSK data rates up to 62Mb/s through use of dual flash A/D’s that sample separate incoming in-phase and quadrature phase components – but the goal of the HAN modem is simplicity and low cost, so the Travertine HAN modem ASIC will modulate and demodulate directly to the HAN frequency range using a single input and a single output signal (single A/D and single D/A). Existing CATV modulator/demodulator chips (such as the Stanford Telecom chips) operate in this manner.

A disadvantage of the simpler QPSK and QAM-4 modulation schemes is the higher baud rates required. With a 14-16Mhz bandwidth allocation, assuming 1 Baud/Hz, the demodulator will need to sample at 56-64 Mhz (4 samples/baud – a typical sampling multiple), and the receiver pipeline will need to run at this clock frequency unless the data is separated into in-phase and quadrature phase components and run through separate receiver pipelines. Although operation at these speeds is facilitated by the simpler 6-bit A/D’s and the advent of .35 and .25 micron technology, max data rates under QPSK are limited by the supportable baud rates. However, these limits to not constrain achieving the stated data rate goal of 20Mb/s.

With the advent of .35 and now .25 micron technology and suitable FIR input and output filters, direct modulation to RF and direct sampling of IF data without the need to mix to and from baseband is possible. The single ended RF I/O of the HAN modem ASIC could then be differentially coupled to the UTP using a suitable BALUN. The Stanford Telecom STEL-1109 is an example of a CATV transmitter chip that modulates directly to RF at carrier frequencies ranging from 5 to 65 Mhz.

Filter Designs

The goal of the Travertine HAN system is to operate, perhaps at some reduced data rate/power level, without the need for any user installed filters or terminators. Filter placement at the TNI can be avoided by lowering the transmit power level to a suitable level, compensating for effects of CPE is more difficult. ***This facilitates a low installation barrier.***

The HAN performance can be improved by the installation of simple passive filters. These filters would be located at the legacy CPE such as POTS phones, and at the Telephone Network Interface (TNI) near the demarcation point between the subscriber loop and the premises wiring. At this time Travertine expects the installation of passive filters in front of the CPE will be required, whereas the TNI filter – which represents the major installation barrier, will be optional.

Figure 9 shows a diagram of the augmented HAN incorporating these filters:

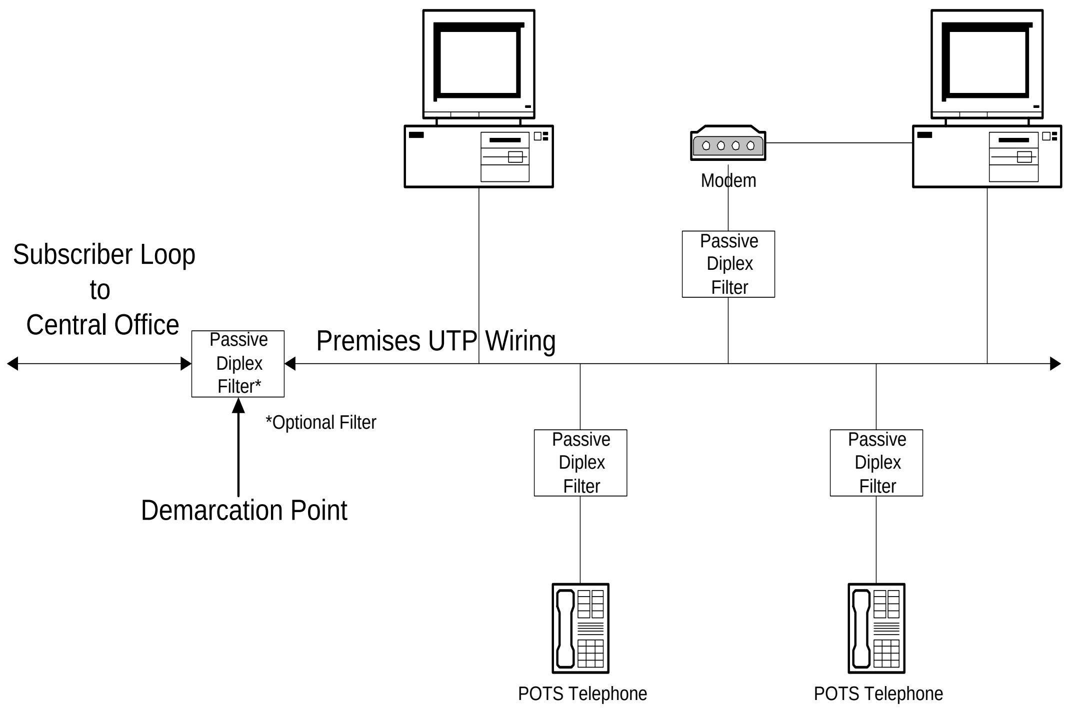

Figure 9: HAN implementation including Diplex filters

In this implementation, passive diplex filters have been installed. All of the user visible filters are designed for easy installation.

There are actually five filters in the system (only two of which are user visible), the diplex filter in the TNI, the diplex filters installed in front of the CPE (such as POTS phones), the HAN modulator output filter, and the HAN demodulator receive filter, and the multiple line support filter(s).

The function of each of the five main filter types is described below:

TNI filter (passive):

a) to keep subscriber loop ingress out of the spectral range of the HAN
b) to keep HAN transmit power off of the subscriber loop to meet Part-68 and/or Part-15 (emissions) requirements

CPE filter (passive):

a) to keep the CPE from putting energy into the HAN spectral range. This would include events such as ringing, trip ringing, pulse dialing, hook switch signaling, etc.

b) to keep the HAN spectral energy from being aliased down into the working frequencies of the CPE
c) to provide a more favorable termination/impedance of the CPE to the HAN in the HAN spectral range. Even if it is to the point of simply reflecting/absorbing transmit energy in the HAN spectral range – this would be preferable over the potentially non-linear impedance/behavior presented by the CPE

Transmit Filter (passive/active):

a) to attenuate HAN transmit power below 1.1Mhz – which would interfere with POTS or xDSL, and to appropriately shape the output pulses to minimize ISI and allow use of matched filters at the receiver.
b) To attenuate HAN transmit power above the HAN passband, primarily for EMI/EMC considerations.

Receive Filter (passive/active):

a) to remove spectral energy below the HAN passband which might be aliased into the working spectral range of the HAN.

Multiple Line Support Filter (passive/active):

a) to effectively "short together" multiple phone lines in the premises into a single network in the HAN spectral range. Unlike the TNI and CPE filters, these filters are designed NOT to pass DC currents so as to allow independent operation of the POTS lines.

The TNI and CPE filters are implemented with passive components, capacitors and wire wound inductors. A likely implementation is as classic ladder filters of the Butterworth/Chebychev type. Metal Oxide Varistors (MOV) may be placed between the TIP and RING lines on both the input and output of the filters to provide surge protection in compliance with Part 68 requirements.

The TNI and CPE filters are of the diplex type. A lowpass filter allows the POTS and xDSL signals to pass through unimpeded, while the highpass section allows termination of the HAN at a known impedance to minimize reflections. Figure 10 shows a conceptual block diagram of the diplex filter used at the TNI:

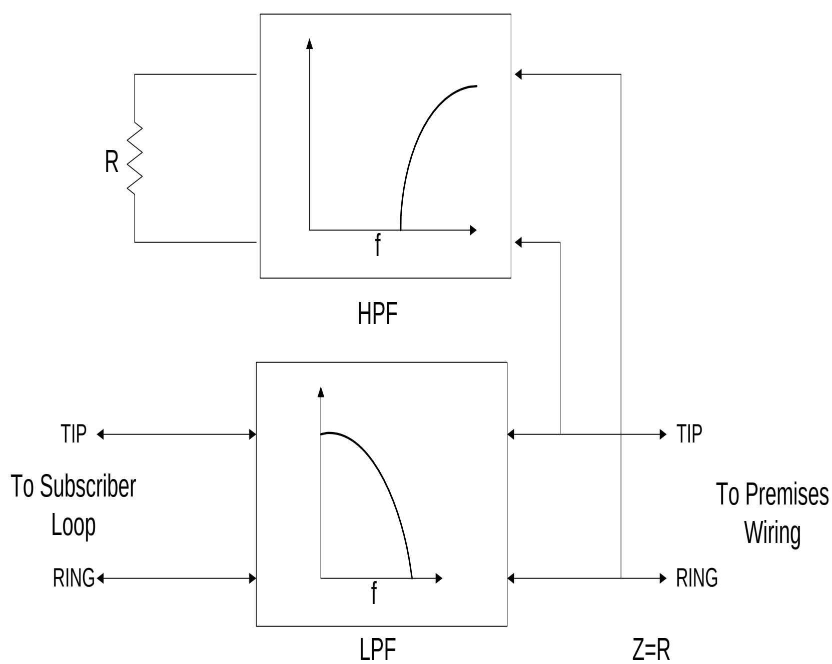

Figure 10: TNI Diplex Filter

The goal of the HPF portion of the diplex filter is to terminate the HAN in its characteristic impedance so as to minimize reflections in the HAN passband. At HAN frequencies the characteristic impedance of UTP is almost purely resistive with a value of about 100 ohms.

The CPE filter consists of a small unit with an RJ-11 jack on one side and a short cable with a female RJ-11 connector on the other. The filter is ideally installed into the wall jack and then the cable to the CPE is plugged into the filter unit.

In order to ease the installation complexity of the diplex filter at the TNI, a clamshell filter unit has been designed such that no cutting and splicing of the UTP wiring is required (no tools are needed). Figure 11 shows a diagram of this filter:

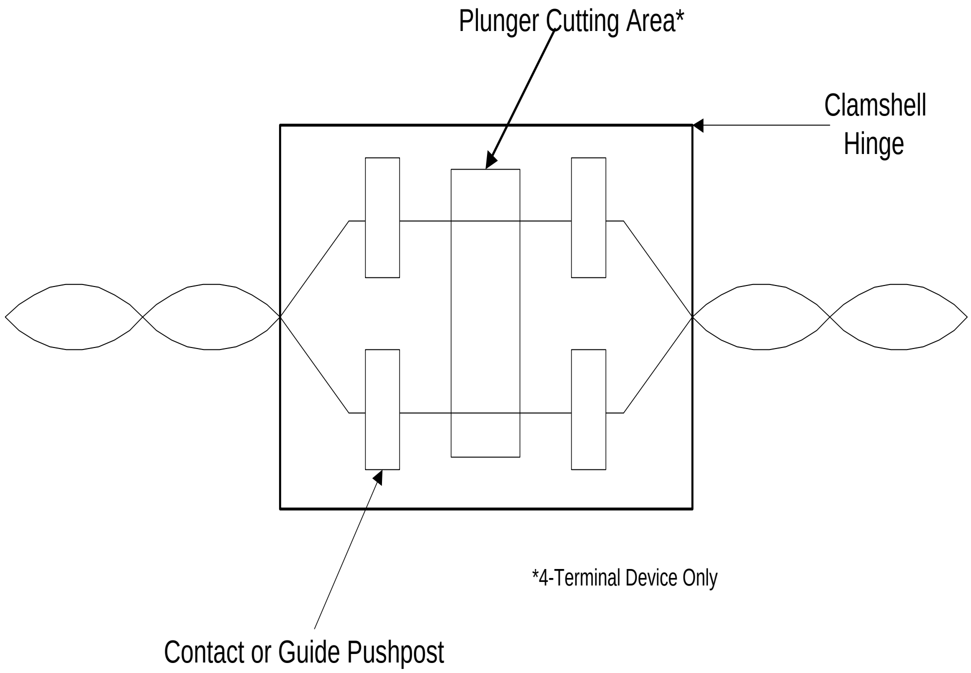

Figure 11: Clamshell TNI Filter

The TIP and RING wires are pushed through pushposts, which can either be guide pushposts or contact pushposts. In the case of contact pusposts, the insulation is cut and a gas tight contact is formed. In one instantiation of the filter, ferrites are used to provide a crude inductance (more of a frequency dependent resistance) and by closing the filter clamshell, the two halves of the ferrites that surround the TIP and RING wires make contact. The ferrite block(s) will typically be grooved to guide the wire through the block. This results in a “zero terminal” device that does not require cutting the TIP and RING wires. In another instantiation, two of the guideposts are contact guideposts, which cut the insulation, allowing for use of a shunting device such as a capacitor between the TIP and RING wires. Because of the crude “inductance” offered by ferrites, and the poor filter characteristics that result, a zero or two terminal filter constructed of ferrites and/or capacitors would only be appropriate for premises that are POTS only.

In Figure 11 above, four contact pushposts are used, and the user depresses a simple cutting plunger, severing the wires and creating a four terminal device, which allows arbitrary passive LC ladder type filters to be constructed. Such a filter would be the likely choice in the case of xDSL usage, where sharper filter skirts are required and hence tuned circuits with the much higher Q that can be offered with conventional wire wound inductors are used.

A conservative filter design allows us to pass the load impedance to the input unimpaired over the POTS and xDSL frequency ranges. This is aided by having adequate transition zones between the upper edge of the ADSL band and the lower edge of the HAN passband.

If ADSL is supported, the "preferred embodiment" for the HAN system filter implementation is one - to use the standard nomenclature – of using distributed (among the CPE) split (into HPF and LPF sections) POTS splitters. This architecture provides the greatest installation flexibility for the user. Figures 12 and 13 diagram such a HAN system. *Note that three different "networks" are hosted by the same PC, and that virtual gateway software is used to interconnect the networks.* See the Virtual Gateway section for more discussion of that implementation.

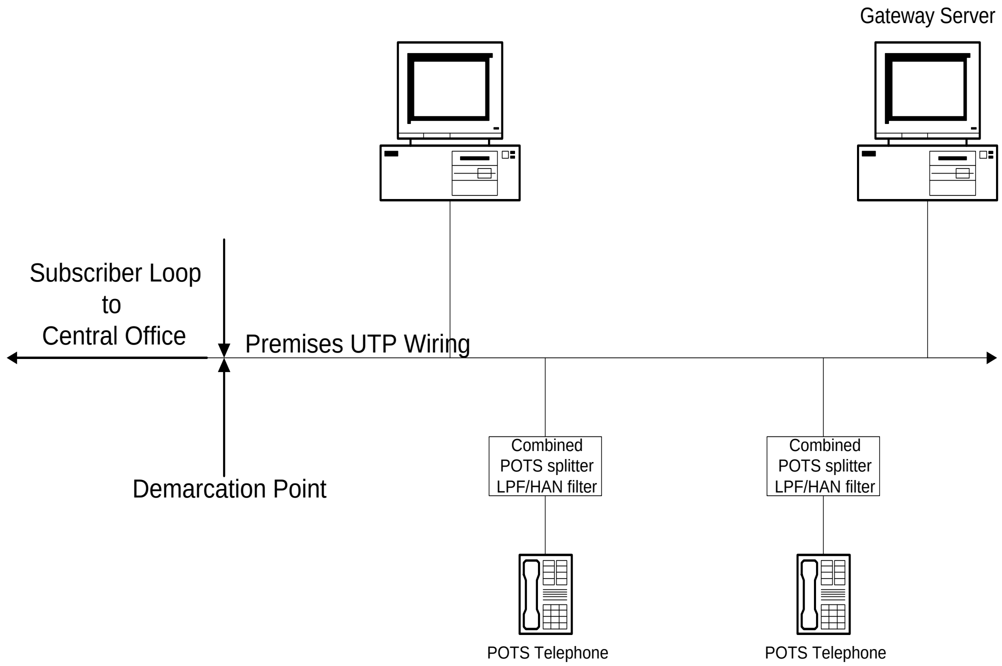

Figure 12: HAN Implementation with ADSL

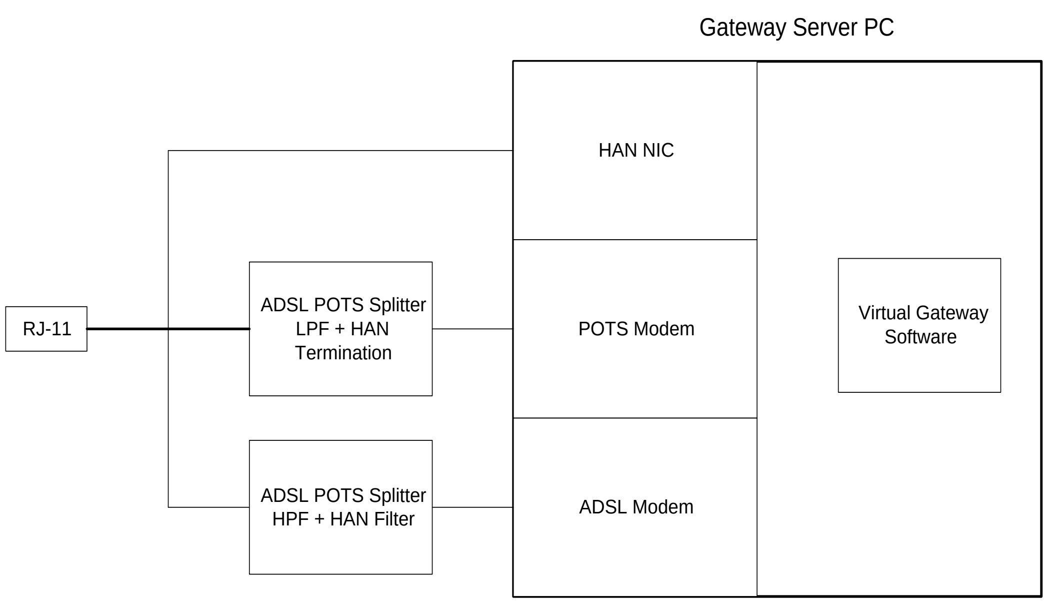

Figure 13: Virtual Gateway Server

HAN Filter Frequency limits

The self resonant frequency of typical wire wound inductors capable of handling the required POTS signaling currents/voltages (without saturating) are usually in the 20 –25 Mhz range. This places an upper limit on practical passive four terminal HAN TNI and CPE filters.

POTS Splitters

Experiences with POTS splitter designs for xDSL applications are instructive in the design of the HAN filters. The return loss spec has proven more difficult to meet than insertion loss for POTS splitters, particularly for Europe. This is principally due to the fact that the impedance presented by the passive filter to the line was not strictly resistive at the frequencies of interest. However the design of the HAN filters is simplified by the fact that the characteristic impedance UTP is almost purely resistive in the HAN passband.

Also, the return loss spec for the US is much easier to meet than the European spec (there are basically no passive European POTS splitters for this reason - they area all designed as active devices).

Side Tone and Echo

There is a relevant article "ADSL and VADSL Splitter Design and Telephony Performance" in the IEEE Journal on Selected Areas in Communications, Dec 95. where the authors state that "the main concern with the effects of a passive filter for analog telephone is its effect on impedances presented to the telephone and to the central office line card. If these impedances are greatly different from the telephone transmission bridge balance impedances they will cause sidetone and echo." "The underlying difficulty of the passive filter approach is the inherently real reference impedance that passive LC filters must have".

Filter DC current requirements

Clause 10.1 (of ANSI T1.413-1995) states that the requirements of T1.413 must be met with DC currents of 0 to 100 ma. However, some old step-by-step switching systems may provide a current as high as 140mA based on 400 ohm battery feed resistors and 56 volt DC supply voltage.

In the Code of Federal Regulations (CFR) 47 part 68 one can find two references to maximum DC currents of 350mA.

For the purposes of HAN filter design, support of a continuous DC current of 350mA is assumed.

Filter Insertion Loss

Probably the critical factor in filter insertion loss considerations, particularly for the TNI filter, is the xDSL minimum received signal level. By the time the xDSL signal enters the customer premises it will be significantly attenuated, hence any pass-band insertion loss will degrade xDSL performance.

From a regulatory standpoint, there is no specified insertion loss except at voice frequencies. Clause 10.2.1.1 (of ANSI T1.413-1995) specifies no more than 1.0 dB insertion loss at 1004 Hz and clause 10.2.1.2 specifies no more than ±1.0 dB attenuation distortion at any frequency between 200 Hz and 3.4 kHz. Because the digital portion of the ADSL is operating under marginal conditions at best, the insertion loss at frequencies of operation (20 kHz to 1.1 MHz) must be minimized. Clause 6.12.1 allows up to 3 dB of passband ripple from 30 kHz to 1.104 MHz for the ATU-C (from 30 kHz to 138 kHz for the ATU-R). Although it may be unrealistic, the average insertion loss at operating frequencies probably should be less than 0.5 dB (while still allowing 3 dB ripple). And since HAN systems may insert two filters in the ADSL path (a filter at the TNI as well as a CPE filter in front of the ADSL modem*), insertion loss will be limited to .25dB, with an even lower value as the design goal.*

The transmission response of double-terminated ladder filters (such as Chebychev/Butterworth) is generally insensitive to changes in the source/load impedance as long as the impedance changes are the same. However, in the case where the load impedance (say at some CPE) is significantly different from the characteristic impedance of the line, the filter insertion loss may vary. This issue is further complicated in the ADSL frequency range by the dependence of the characteristic impedance of the line on frequency. This issue deserves further study.

There is some risk that the HAN filter(s) (through insertion loss) and/or the HAN transmitters (by putting energy in the ADSL frequency range) may degrade ADSL performance – however given the low expected penetration of ADSL into the customer premises, and the principal focus of the HAN on intra-home PC networking, this risk is considered acceptable.

Filter Return Loss

Clause 10.3.1 ((of ANSI T1.413-1995) requires 10 dB or better return loss relative to 100 ohms from 30 kHz to 1.1 MHz. The 100 ohm impedance requirement recognizes that the characteristic impedance of twisted pairs converges to 100 ohms resistive at high frequencies. As for the POTS interface, the required impedance is 900 ohms at the central office end and 600 ohms at the CPE end. The required return loss, relative to the respective impedances, is remarkably low. Clause 10.2.1.4 and Table 29 require between 5 and 8 dB return loss depending on the part of the voice frequency band being measured. T1.413 breaks the voice frequency band into three regions where singing return loss and echo return loss are to be measured. These band regions are standardized in IEEE Std. 743-1984

Filter Group Delay

How much Envelope Delay Distortion (EDD) can the filter add to signals incoming from the subscriber loop so as not to interfere with voiceband or xDSL transmissions? Clauses 6.12.1 and 7.12.1 require no more than 50 microseconds over the passband of 30 kHz to 1.104 MHz for the ATU-C and from 30 kHz to 138 kHz for the ATU-R. Clause 10.2.1.3 require no more than 200 microseconds between 600 Hz and 3.2 kHz for the POTS interface.

Something to consider though is that the subscriber loop presents a constant group delay because the phase change is relatively linear with frequency, whereas the filter(s) group delay characteristic can be highly non-linear. And while the specifications of ADSL receivers may lead one to believe they can handle the group delay, it is critical from a marketing perspective that the HAN not compromise even marginal ADSL performance.

Required Filter Attenuation

TNI Filter

How much attenuation must the TNI filter provide at the HAN passband so as to be compliant with Part 68 on the subscriber loop? Part 68 specifications state that the maximum is -15dBV for 135 ohm metallic impedance for frequencies from 270Khz to 6 Mhz. This corresponds to a voltage level of .18V for a 100 ohm system.). Then from a Part 68 standpoint, the amount of attenuation necessary from 1.1Mhz to 6 Mhz is minimal for power transmitted from the premises wiring onto the subscriber loop (see Allowed Transmit Power section for a more thorough discussion).

Another use of the TNI filter is to attenuate subscriber loop energy in the HAN spectral range (to keep subscriber loop noise from interfering with the HAN). This is important from keeping subscriber loop impulse noise and ISDN/xDSL generated NEXT from leaking into the premises HAN. This attenuation value is still TBD. A significant amount of immunity to NEXT is gained because only downstream ADSL signals come close to the HAN passband. And since these signals must originate at the CO, they will be heavily attenuated by the time they enter the premises. The much higher attenuation of UTP in the HAN spectral range in general provides noise immunity from subscriber loop ingress.

HAN Transmit Filter

For HAN transmitters to not interfere with incoming xDSL signals from the subscriber loop is believed to be the most stringent of the HAN filter requirements, and that 90dB of attenuation may be needed to reduce the HAN transmit energy to essentially thermal/background noise levels in the stop-band.

Filter Component Protection/Selection

Capacitors must be voltage rated to 300V P-P. Inductors must be current rated to 350mA. MOVs are used between TIP and RING on the filter input and output, and are specified to break down at 250V.

Medium Access Protocol (MAC)

The MAC and Network Management layers of the HAN are just beginning development. The MAC layer in particular will require support at the physical layer level, so fleshing out the MAC is a high priority. However some issues have been explored:

Distributed vs Centralized

Protocol implementations can be simplified if there is a bus master to deal with arbitration, prioritization, and centralized timing issues. Travertine may make the assumption of a centralized intelligent resource in order to simplify the cost of the slave nodes. While a completely plug-n-play model is the goal, this will be relaxed if the network implementation can be significantly simplified.

Service Types

The network supports peer to peer half-duplex connections, and broadcast communication. Situations where an unknown number of stations can be transmitting, or where the identity of the transmitter is not known cannot be supported because the equalization would not be appropriate. Use of CDMA for control slots to handle service requests and other housekeeping communications by multiple simultaneous requestors can probably be supported if the number of transmitters is always known to the target receiver and the receiver's equalizer can be appropriately trained.

Initialization

In order to (auto) initialize properly, the network must be operated in a domain that is free of equalization requirements. For QPSK this can be accomplished by operating at baud rates in the 500Kbaud range, where the baud time is sufficiently long to allow the reflections to be naturally attenuated below the QPSK SNR threshold.

Another possibility strategy is to leave the spectrum allocation or baud rate alone, but instead to employ direct sequence spread spectrum modulation (CDMA). Spreading the effectively narrower band information over the available spectrum mitigates reflection effects. The chip rate (1/code clock) would be the same as the normal system baud rate under QPSK, but the overall data rate would be at least an order of magnitude less. The effective data rate is similar to a system with a lower baud rate, as above, but the effective baud/chip rate and bandwidth allocation can stay the same, thus bypassing the dwell time that would inevitably occur in systems that change buad rates, constellation sizes, and/or bandwidth allocation on the fly. CDMA systems typically employ QPSK or BPSK modulation for the chips.

Training is an O(n) process, where each node transmits and every receiver simultaneously trains its equalizer to that transmitter. For every peer to peer connection, there are two sets of coefficients, one for each transmit direction.

Network Control

Situations where either the baud rate or the constellation size changes dynamically to support control channels and/or to support different capability devices are not desirable because of the system dwell time likely incurred when switching modes.

Carrier Sense & Collision Detection

The baseband access protocol CSMA/CD is used for Differential Manchester encoded Ethernet systems. A key component of the protocol is the ability to quickly sense a carrier and detect a collision by examining the line voltage level. This baseband collision detection and carrier sense approach is not a viable one for a RF passband system such as that proposed for the HAN.

Use of received signal power is difficult, but not impossible, as a collision detection method for the HAN. This is because of the greatly varying strengths of possible transmitters on the network as seen by any receiver in the highly attenuative UTP environment of the HAN in the HAN spectral range. A distant colliding transmitter could easily be 10+dB down from the primary transmitter and might only raise the total received signal power by 1dB, which is expensive to reliably detect. A more reliable indication would be an invalid block code/checksum or change in the BER, but this might be valuable for other uses (see below). Existing passband Received Signal Strength Indicator (RSSI) implementations typically measure received signal power over many symbols and are typically only +/- 3dB accurate.

Carrier sense is also difficult in a passband system utilizing QPSK modulation. Neither QPSK nor QAM does not have a strong spectral component at the carrier frequency (no residual or pilot carrier). Alternative methods are employed for carrier syncronization as part of timing recovery, but this is a more involved and time consuming process. Differential encoding is typically used to resolve any ambiguities in the recovered carrier.

Detecting changes in block codes/checksums or the overall BER work best for passband systems that are not dynamically reconfigurable, or subject to dynamic changes in their impedance/transmission environment. However, the HAN network may experience impedance or network configuration changes during operation, and using changes in the bit error rate to detect transmission collisions renders it difficult to use this similar techniques to detect changes in the system configuration/impedance. This coupled with the known inefficiencies of the Ethernet CSMA/CD protocol under heavy loading are leading Travertine to closely examine TDMA (token) protocols for the MAC layer. Although more expensive to implement, TDMA protocols are more efficient under heavy loading conditions, provide traffic prioritization to providing QOS for isochronous data, and would not only allows changes in bit error rate to signal a network configuration/impedance change, but may also open the possibility for other, more active, techniques for examining the network configuration (discussed below).

As with the case of centralized vs distributed control however, if the HAN impedance/spectral characteristics prove stable, with the CPE filters installed, to anything other than the attachment of a new device to the network – then user of a BER change to signal collision might be used if the system implementation can be considerably simplified – at the possible expense of plug-n-play robustness.

Detecting Network Changes

The network characteristics will change whenever a new device is attached and/or any length of cable is attached to the network. If CPE filters are not installed, just taking the phone off hook can significantly change the network characteristics – necessitating retraining of equalizer (and other) coefficients.

The two methods currently under consideration for use in dynamically detecting changes to the HAN configuration/impedance are:

a) changes in the uncorrected bit error rate, and
b) dynamically measuring the network using active TDR techniques

Once a change has been detected, the system can either:

a) request a retransmission,
b) request a new set of stored equalizer filter and other parameter/coefficient data be loaded at the demodulator (and potentially at the modulator as well)
c) request that the system retrain

Changes in Uncorrected Bit Error Rate

The FEC logic of the HAN demodulator includes indicators of bit error rate (BER) and/or provides an interrupt type signal when a preset threshold is exceeded.

Active Time Domain Reflectometry (TDR)

The flash A/D sample rates of high baud rate QPSK systems are nearly fast enough to support meaningful TDR measurements of the premises network. Furthermore, for discriminating measurements to be taken, shaped pulses on the order of 10ns or less must be generated, and a 10ns pulse period is very close to the baud period of the HAN under QPSK. One proposal under examination by Travertine for the integrated HAN modem ASIC is explicit support for this capability. The system would, on a heartbeat basis, perform a TDR measurement on the network looking for changes. If a change is noted, more deterministic responses can then be taken, such as selecting exactly the right stored coefficients/parameters at the transmitter and receiver.

Cells vs Packets, or ATM vs TCP/IP

The PC and consumer electronics manufacturers prefer TCP/IP systems, while the Telcos desire to deliver ATM to the desktop. Travertine is relatively agnostic regarding this issue (currently). But while TCP/IP can be delivered over the layer II ATM (the LANE proposal), given the low penetration rate of WAN to the home via the subscriber loop, and Travertine's focus on the PC market as its initial market, supporting TCP/IP ("virtual ethernet") makes the most sense. If the Telcos eventually do deliver ATM to the premises via xDSL, then the protocol conversion can be done locally within the premises (via the Virtual Gateway). This makes more sense than doing the conversion at the central office as it presents no barrier to the Telcos should they insist on ATM to the premises, and it distributes the protocol conversion compute load (and power consumption) among the premises.

Noise Sources

Possible noise sources the HAN must contend with include:

a) power line coupling
b) impulse noise from dimmer switches, flourescent lights, electric motors, and from events such as ring signals, or receivers being taken on and off hook
c) CW noise from radio transmitters
d) Crosstalk from other lines in the house
e) Subscriber loop NEXT from adjacent lines in the cable
f) Subscriber loop impulse noise

Thermal/Background Noise

A background Noise Power Spectral Density (NPSD) of -140dBm/Hz is a conservative value for use in link budget and other calculations. It is a value frequently cited in the literature.

NEXT

NEXT sources include other transmitters in the home which use other twisted pairs in the same biquad or in-house wiring bundle, and include crosstalk from adjacent pairs in the subscriber loop cable. The latter can be attenuated by use of a TNI filter, and as previously mentioned, we are aided here in the case of ADSL because ADSL upstream spectral allocation is far from the HAN passband.

Impulse Noise

Impulse noise is generally defined as a voltage increase of 12dB or more above the background (rms) noise lasting 10ms or less.

Impulse noise (from ringing and trip ringing in adjacent pairs, as well as from sources such as lightening), can have "destructive" amplitudes as long as 1ms. This is the primary impetus for Reed Solomon FEC (forward error correction) in the ADSL standard. Interleaving can also be used to further correct error bursts of 500 microseconds, but at the expense of latency (17ms for 500 microsecond bursts).

Forward error correction techniques such as Reed-Solomon are effective in combating impulse noise, but as bit rates increase so does the need for interleaving to temporally spread the impulse events over the data stream.

CW Noise

CW sources such HAM Radio radiate in the HAN band, but are not expected to present a serious ingress source issue except in rare cases.

Error Correction

The HAN employs three forms of error correction:

1) Reed Solomon FEC (to correct burst errors)
2) Interleaving, to obtain better Block Error (BLER) performance. This adds significant latency – on the order of 40x the maximum length correctable impulse.
3) Trellis/Viterbi convolutional encoding, which provides protection against Gaussian noise such as crosstalk and background white noise.

Because of the heavy latency cost imposed by interleaving, and the memory cost associated with its implementation, minimization of large impulse noise sources is of high value, and is one motivation for isolation of the subscriber loop (via the TNI filter) and its impulse sources such as lightning.

Equalization

Aggressive equalization is critical to the performance of the HAN. In the HAN environment, channels are asymmetric from an electrical standpoint, and any given peer to peer channel will require different coefficients for each direction.

Training occurs during network initialization, during network idle times, and when specified by a receiver. Training is typically performed by the transmission of a known bi-level training sequence. The receiver correlates the training sequence and updates the equalizer filter coefficients using an adaptation algorithm such as Least Mean Squared (LMS).

Typical CATV chip equalizers can take 200ms to converge the equalizer coefficients (from a cold start). This latency is intolerable if training had to occur every time a peer to peer channel was established. Hence, fast loading of pre-trained equalizer coefficients (and other transmitter/receiver parameters) is a key component of the HAN system implementation. Since existing off-the-shelf technology does not support on-chip storage of equalizer coefficients and other parameters, a fast host interface for access to the on-chip equalizer coefficients is a key requirement. It is not known at this time how much correlation there may be between equalizer coefficients for various peer to peer channels on the same HAN – if significant correlation exists the training time could be less.

Equalizer Tap Length and DSP Compute Equivalent Analysis

A relevant issue is the computational requirements to implement the HAN physical layer on a DSP. Lets just take the adaptive equalizer and perform a "back of the envelope" analysis. For acceptable QPSK reception, the noise power due to multipath/echo needs to be ~15dB down from the signal power. For QAM-16 the noise power needs to be ~25 dB down. Knowing the transmission frequency, we can calculate how far the signal needs to travel before it is naturally attenuated to the desired level. Once this is known, then for a give baud rate the required number of filter taps can be calculated.

The most difficult frequency from an equalization standpoint (# of taps) is the lowest one in the band (our current lower band edge target is ~4Mhz), where there is the least amount of attenuation. The attenuation equation in the ANSI-568A standard gives a cable length of 850ft to achieve 15dB down @ 4Mhz. This matches very closely what I found in my own house, which was about 8dB down for 450ft @ 4Mhz. So for 15dB roughly 850/.64 (0.64 ft/ns propagation speed) or 1.3 microseconds worth of equalization is needed.

The DVB standard - to which the CATV demodulators are designed to - specifies 1.1 microseconds of equalization at 7Mbaud. Existing chips typically equalize from 1.1 to 2 microseconds. Note that one can take advantage of the dependence of the natural attenuation on frequency to decrease the equalizer length by increasing the frequency of the lower edge of the band. But you can only do this so far until the upper band edge is attenuated so heavily that it disappears into the background noise.

Note also that increasing the constellation size (say to QAM-16), while halving the spectrum required for the same bit rate, increases the length of the equalizer if the frequency of the lower edge of the band stays the same. This is because of the increase in SNR required (10dB). Hence, if the band size and center frequency remain the same, the bit rate cannot be increased by increasing the constellation size unless the equalizer tap length is increased - this is independent of signal power! For a house with 500ft max peer-to-peer distance, 10dB of increased attenuation corresponds to moving the lower edge of the band up in frequency by about 10Mhz (from 4Mhz to 14Mhz).

The LSI Logic CATV chip uses a DFE equalizer with a total of 32 coefficients to achieve 1.1us of equalization at 7Mbaud, while the VLSI Technology chip supports either DFE or transversal equalizers with 52 total coefficients and is capable of canceling echos out to (at least) 1.6 us. 1.1us is equivalent to 8 baud times at 7Mbaud, and since all the CATV parts use flash A/D's that sample at 4x, this works out to 32 taps. The equalizer is clocked at 1/2 the sampling frequency, or twice the baud rate. Equalizers used for ISDN are typically clocked at twice the baud rate.

Let us assume 1.5us worth of equalization for a 10Mbaud HAN under QPSK. For a 10Mbaud QPSK signal - a baud rate that provides a net data rate in the mid to high teens after error correction overhead is accounted for - this is equivalent to 15 baud times. So at a 4x sample rate, this gives a 60 tap filter. The minimum 8 x 8 MAC rate is therefore 60 * 20Mhz = 1.2G MACs/s. This tells us that the equalization, filtering, and FEC compute requirements of high baud rate QPSK systems are beyond a cost effective software only (current DSP technology) implementation. The advent of .35 micron, and particularly .25 micron feature sizes make high baud span equalizer implementations tractable.

The HAN QPSK modem is a good candidate for dedicated silicon because the compute problems are highly systematic, the encode/decode is simple (not much time spent in shell mapping and demapping), and the algorithms are compute intensive using non-standard multiplies (the QPSK data consists of only a few bits – allowing considerable silicon area to be saved by using specialized multiplier designs).

BALUN

BALUN transformers will be required to match the single ended ground referenced I/O of the HAN modem ASIC or modulator/demodulator pair to the differential twisted pair premises wiring. Wideband BALUNs with the necessary return loss, longitudinal balance, common mode rejection, and output signal balance will be required. Suitable devices have not yet been specifically identified for the HAN, but BALUNs for ISDN U interfaces, HDSL front ends, and T1 interfaces are readily available from companies such as Pulse Engineering at a cost of $2-3 in volume. As previously mentioned, proper resistive termination structures can significantly reduced conversion of metallic to common mode signals and thereby reduce emissions.

The isolation requirements of Part-68 compliance also make some form of transformer or switched capacitor (such as that employed by Krypton) isolation necessary.

The AFE/UTP interface section of the HAN is considerably simpler than a standard modem DAA as there are no ring detect, off-hook, pulse dial, caller-ID, line current hold or other functions that need to be implemented.

Multiple Line Support

In keeping with the mass market consumer spirit of the HAN, there needs to be a method of handling the common situation where multiple POTS phone lines are used in the home. A typical example is where a second line in the biquad is activated for use by teenagers, and Figure 14 shows such an installation:

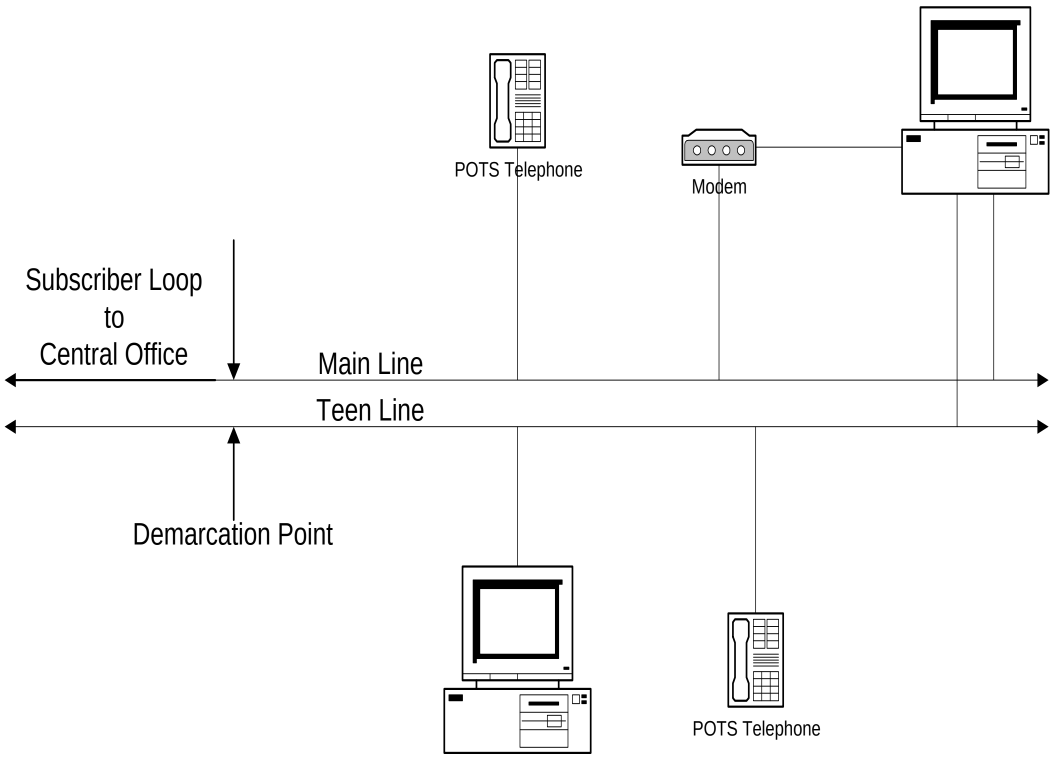

Figure 14: Multiple Lines in the Premises

There are a variety of possibilities as to how the lines are made available in the home: some wall jacks may be dedicated to one or the other line, some wall jacks may be wired to both lines, and in some locations dual jacks are available, where each line is brought out to a separate jack.. We would like the user to be able to plug into any wall jack and utilize the HAN without much thought to household wiring topology. All that is required is that somewhere in the network, a HAN client is used to “short together” the two networks in the HAN spectral domain, leaving independent POTS and xDSL service on the other two lines unaffected. This is a preferred solution over alternatives, such as requiring the user to install multiple NIC cards and have the PC perform bridging functions between the networks. This is shown above at the computer in the upper right of the diagram where a location with dual jacks is used to tie the lines together. Depending on where the lines are tied together, the maximum effective peer-to-peer distance can increase in such a HAN topology. This has implications for signal attenuation in the link budget analysis.

Figure 15 shows a diagram of a general implementation of a solution to “shorting together” the two lines. Note that the two lines can present themselves to the hardware as either a single RJ-11 jack (using the TIP1/RING1 and TIP2/RING2 pairs) or via two separate jacks.

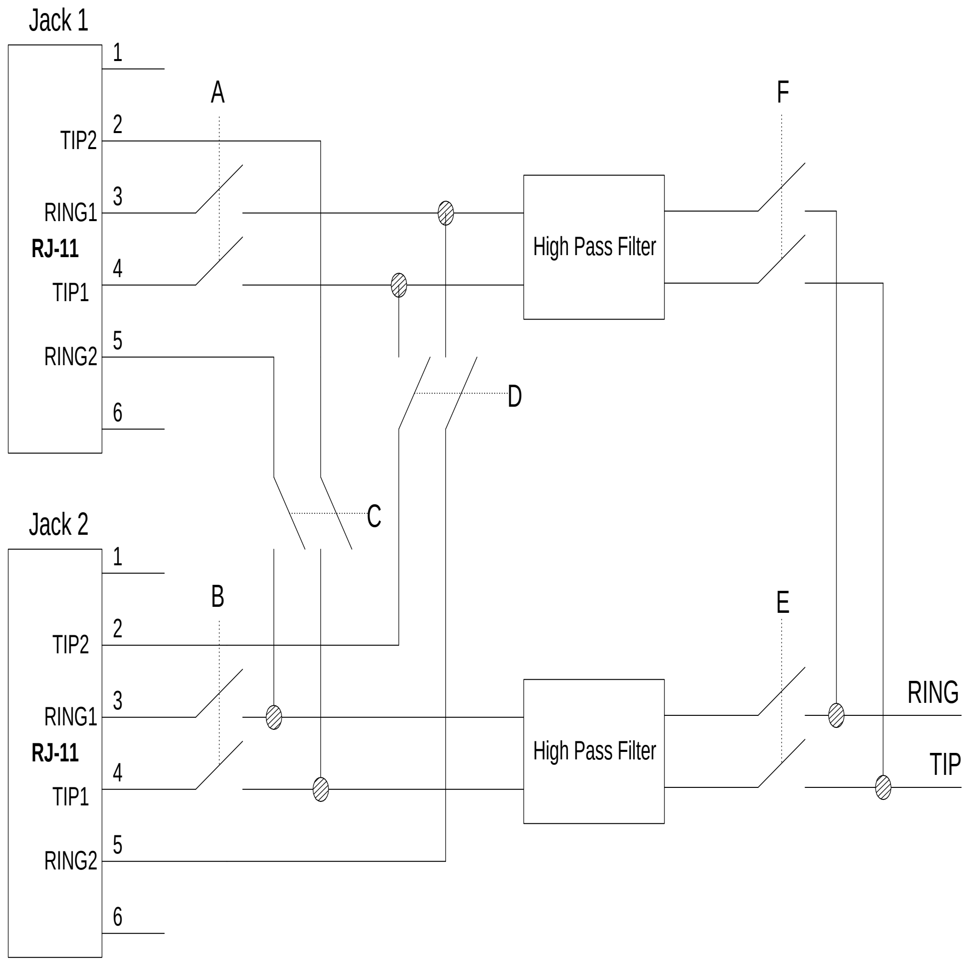

Figure 15: Joining multiple lines

Note that the high pass filters used in this circuit must be designed to block DC currents so as not to interfere with the independent operation of the POTS lines. If passive LC ladder filters are used, the filters would lead with capacitors and use inductors as the shunt elements.

Virtual Gateway

The Virtual Gateway concept refers the Gateway/Bridge functionality that the HAN can offer to multiple distinct networks sharing the same electrically contiguous communications medium. The most likely disparate networks to be supported are POTS (acting as a WAN), and the HAN (acting as the LAN). By electrically contiguous medium we mean one where DC current can pass between any two communicating devices without any active device(s) being inserted in the communications medium between the communicating devices. The passage of DC current between the Telco Central Office and the CPE is essential for POTS lifeline service.

By virtue of its passband/FDM implementation, The HAN supports implementation of Gateway functionality without the need for an active physical device that terminates the subscriber loop at the premises. A software only functional equivalent, running on a PC, can be performed if the PC is equipped with a POTS modem, and/or an ADSL NIC, and a HAN NIC. The PC can forward data from the subscriber loop WAN not by physically transferring the data from the subscriber loop to the premises wiring through an active device, but by "hopping" the data up to the HAN spectral range and retransmitting. The implementation of the "Virtual Gateway" can be either centralized, partially distributed, or fully distributed – and these versions are discussed below:

Centralized

A single Gateway Server attached to the HAN that can send/receive on either the WAN or the HAN. In this case the other HAN devices (which for descriptive purposes we will label as "HAN Clients") need only send/receive in the HAN spectral domain, and the Gateway Server will "hop" network traffic between the WAN to the HAN spectral domains. This implementation results in the lowest complexity HAN Client devices, however it is wasteful of HAN bandwidth because any WAN data destined for HAN Clients has to be retransmitted on the HAN, and any HAN Client data bound for the WAN has to first appear on the HAN.

Partially Distributed

A Gateway Server attached to the HAN that can send/receive on the HAN, but may perhaps be the only transmitter on the WAN that is "HAN capable". The HAN Client devices could send/receive on the HAN, but perhaps only receive on the WAN (typically made accessable via ADSL). In this scenario, HAN Clients can receive WAN data directly, but must "hop" data destined for the WAN through the Gateway Server. This increases the complexity of the HAN Client devices, however, it creates an ideal situation for the receipt of broadcast WAN data, such as VOD (Video On Demand) while preserving as much HAN bandwidth as possible. Note that both the Gateway Server and the HAN Clients implement Gateway/Bridge functionality.

Fully Distributed

In this implementation, there is no centralization of any of the Gateway/Bridge functionality, and any HAN Client device can send/receive on either the HAN or the WAN. This results in the most flexibility and system robustness, and the expense of HAN Client implementation complexity.

Link Budget Analysis

As a sanity check we perform a back-of-the-envelope link budget analysis to get a feel for the system limits for a system with a passband from 4 to 18Mhz. For this first order analysis we ignore noise contributions such as crosstalk, CW noise sources, and power line coupling. Because attenuation losses and multipath effects so dominate this environment, for now we also ignore BALUN and filter insertion losses and other similar effects.

We choose a wide-band background white noise NPSD of –140dBm/Hz (which is prevalent in the literature) as a conservative number for UTP applications (a pure thermal NPSD at room temperature would be –174 dBm/Hz). The total noise power contribution of a 14Mhz HAN band is then -68.5dBm. Let us assume we are transmitting at 10dBm (1Vp-p). The signal is usually launched into a tap that splits into two directions once the main line is reached, incurring an immediate 3dB loss (worst case – particularly at the high end of the band).

If QPSK is used then a 15dB SNR is required at the receiver for an acceptable error rate, and we choose to add 6dB of margin. If we ignore the effects of other bridged taps, this gives an available UTP attenuation budget of 54.5dB – which at the upper edge of the HAN passband (18Mhz) allows peer-to-peer wire lengths of approximately 1250 ft @ 18Mhz if the ANSI standard CAT-3 attenuation equation is used.

The presence of bridged taps increases the overall slope of attenuation vs frequency. If we provide at additional 10dB of margin to account for bridged taps, then the resulting 44.5 dB attenuation budget allows peer-to-peer wire lengths of approximately 1000 ft @ 18Mhz.

If we choose a transmit power which is guaranteed to be Part-68 compliant (-5 dBm or .18V) without the presence of a TNI filter, then an additional 15dB of attenuation budget is lost, which reduces the max peer-to-peer wire length to approx. 700 ft @ 18Mhz.

Note that the pronounced "tilt" in the received signal margin (due to the frequency dependent attenuation) over the entire frequency band means that the "average SNR/baud" over the entire band is actually ~13dB higher than the margin at the highest frequency (18Mhz). The exact benefit of the increase in SNR in the lower frequencies of the band is difficult to calculate analytically, but intuitively one can see that it should benefit the overall channel SNR.

If QAM-16 is used instead of QPSK, then the allocated spectrum can be reduced by a factor of 2, which provides an additional 3dB to the budget from the reduction in background white noise power, but the SNR at the receiver must be increased 10dB to 25dB. Hence for the 10dBm transmit power case, the attenuation budget decreases by 7dB over QPSK, and the allowable peer-to-peer wire length given the resulting 37.5dB attenuation budget is 1200ft @ 11Mhz - an *increase* of ~20% over QPSK. But note that for QAM-16, the smaller band translates to a reduction in the “average SNR/baud” to ~7dB due the “tilt” effect noted above.

However, because of the many benefits of QPSK outlined in the “Modulation Scheme” section, QPSK remains the modulation scheme of choice for the initial Travertine HAN. If data throughput rates much in excess of 20Mb/s are required in future, the HAN may migrate to QAM-16 or QAM-64.

In summary, this link budget analysis indicates there appears to be more than sufficient channel SNR margin to meet the goals of the Travertine HAN, even if the allowed transmit power is limited to a Part-68 compliant -5dBm.

Review of Existing Technologies

RF LANs

Clearly an expensive robust high bandwidth RF LAN that adequately addressed security and interference concerns would be the holy grail of home networks. There are products available that utilize the spectrum allocated at the 2.4Ghz and 5.8Ghz bands.

The 2.4Ghz band is limited to about 2Mb/s due to spectrum allocation, but inexpensive all CMOS transceivers will likely soon be available. If the robustness, interference, and security issues can be solved, this technology could be a competitor to the HAN for the lower bandwidth service needs.

The 5.8Ghz band has the spectrum available to provide 20Mb/s and higher service, but faces other difficulties because microwaves at this short wavelength are easily reflected (off of glass, for instance), so that multipath becomes a significant systems issue. To overcome this, antenna diversity and circular polarization are used to discriminate between direct and reflected signals. In fact, a leading existing system actually depends on multipath to recover the signal, and performs poorly in environments that are low in reflecting materials (such as the home). The cost of 5.8Ghz systems is not expected to drop below several hundred dollars per node in the near future because of antenna costs and GAS front ends that are used.

An issue systems in both these spectral domains face, but most particularly the 5.8Ghz band, is radiation patterns and system antenna design. In the 5.8Ghz range a system with antennas designed for a planar office environment may not perform well in a multistory dwelling.

ADSL Modems

ADSL modem chipsets are becoming relatively common, with Motorola, Alcatel, Amati, Globespan, Westell, Pairgain, Aware, Broadcom, TI, Lucent, HP, and Ameritech providing either technology or chipsets. The two prevalent technologies are CAP (Carrierless Amplitude/Phase Modulation) and DMT (Discrete Multitone Modulation). CAP is characterized as “baseband QAM”, and DMT is a sophisticated multi-carrier technology. CAP’s genesis is from Lucent, and Amati would have one believe that they developed DMT, although DMT has actually been around for some time, having originated in the military.

Motorola has delayed the copper gold chipset, and there is confusion in the market ragarding adoption of either DMT or CAP as the defacto standard. These problems are delaying adoption of ADSL, and are likely to encourage the development and adoption of alternate (non-subscriber loop) broadband delivery mechanisms to the home.

Amati holds a number of patents on improvements to DMT (which we understand provide only a 10-15% performance improvement, at considerable compute overhead). Amati was successful in getting the ANSI T1E1.4 working group to adopt DMT as the modulation method for ADSL standard T1.413 in

1993. Although DMT is in theory highly adaptive, both from the standpoint of avoiding spectral areas of ingress, and in rate adaptability. It is however quite complex to implement, has a high peak-to-average power ratio (PAR), is computationally expensive (using IFFTs and FFTs to "channelize" the data), and suffers from long latency due to the large symbol size *(DMT typically performs IFFT on blocks of up to 1K bits)*. The compute requirements necessary to dynamically measure the channel capacity of each of the 256 discrete 4Khz sub-bands is considerable. The symbol latency without error correction overhead is in the 2ms range for DMT. Because of the large symbol sizes, DMT is more susceptible to long burst errors due to burst noise events, so proportionally greater (higher overhead) error correction schemes (interleaving) are required for DMT implementations. The high PAR of DMT (about 17dB higher than CAP), has meant in practice that clipping of peak signals is employed, which requires error correction for recovery. The high PAR of DMT also translates to increased dynamic range requirements and power dissipation for the DMT Analog Front End (AFE). DMT has not yet been chosen as a standard modulation method for any *symmetric* subscriber line standards.

CAP has been on the scene longer, and is a simpler method to implement. The main implementation complexity in CAP is found in the adaptive equalizer employed, and is in part due to the higher baud rates used in CAP. CAP is a standard line encoding scheme for both HDSL and SDSL. Globespan is currently the largest proponent of CAP.

ADSL modems incorporate a variety of technologies that are applicable to the HAN, but these solutions are targeted toward solving the bandwidth problem on the subscriber loop, which is a very different domain from that found in home premises wiring. The fierce attenuation of the 12 km CSA means that there is only 1 to 1.5Mhz of practical spectrum available to work with – as opposed to perhaps 20 Mhz of available spectrum in the home. Because of the physical implementation of subscriber loop UTP into cable bundles, subscriber loop channel performance is typically NEXT dominated from adjacent twisted pairs running xDSL or ISDN. This is not the case in the home environment. In order to support full duplex communications, ADSL modems typically employ echo cancellation or FDM to allow the upstream and downstream channels to share the same precious spectrum. Full duplex communication is not necessary to support the networking requirements of the HAN, but there is sufficient spectrum such that simple FDM can be employed (rather than echo cancellation) if full duplex communication were desired.

Satellite DSS/DVB Receivers

Satellite receivers such as those found in DSS set-top boxes at first glance seem nearly ideal for use as initial off-the-shelf HAN demodulators. They conform to the DVB standard and are optimized for BPSK/QPSK reception. They possess high speed 6-bit flash A/Ds (up to 90 Mhz), matched filtering, carrier recovery logic, clock recovery logic, AGC logic and extensive FEC logic; including variable code size Reed-Solomon block decoding, variable rate Viterbi (maximum likelihood) convolutional decoding, de-interleaving and descrambling. They are capable of QPSK demodulation rates of 45Mbaud (90 Mb/s for QPSK), and these high data rates are an indicator of the robustness of QPSK.

They however are missing one crucial element which makes them unsuitable for use with the HAN – an adaptive equalizer. Satellite modems are designed for direct line of sight channels that are virtually multipath free – a situation distinctly unlike that found in premises UTP wiring environments.

Satellite demodulators do however, make a good candidate as base technology from which to develop the integrated HAN modem ASIC.

Cable Modems

Cable modems conforming the MCNS/DAVIC/IEEE 802.14 provide the closest match to the HAN requirements for available off-the-shelf high volume technology. Their key feature is the inclusion of an adaptive equalizer which, depending on the product, provides from 1.1 to 2.0 microseconds of equalization at the maximum baud rate range of 7 Mbaud. However, cable modem chips include functionality not needed by the current HAN design center, such as:

a) support for higher QAM constellation sizes, particularly 64 and 256 QAM
b) more costly 8-10 bit A/Ds in order to support the higher QAM constellations

They are also lacking in some desireable features, such as:

a) convolutional coding (Viterbi)
b) variable code size Reed Solomon coding
c) explicit support for rapid switching from one transmitter to another
d) local storage for previously trained adaptive equalizer coefficients
e) local storage for additional sets of FIR filter coefficients (transmit filter, and receive/matching filter)
f) local storage for previously trained carrier acquisition, clock recovery, and AGC control parameters.
g) more extensive equalization with potentially finer tap weights

Integrated CATV subscriber modems are not designed to be transceivers, this is because the upstream and downstream requirements are not symmetric. This presents challenges when contemplating use of the integrated single chip modulator/demodulators.

Existing Technology Choice Rationale

Table 1 below presents a survey of likely off-the-shelf technology to be used for introduction of Travertine's first products. They consist of CATV and satellite modulators and demodulators. From a practical perspective, the satellite demodulators must be immediately dismissed because of their lack of adaptive equalization, nonetheless the feature set they provide is instructive as it may make a good base from which to develop the integrated HAN modem ASIC.

The column labels are as follows: Board/Chip, QPSK/QAM-4 support, QAM-16 support, Equalizer support, Data rate, Scrambler, Differential Encoding, Reed Solomon FEC, Convolution coder/Viterbi Decoder, Interleaving support, programmable Nyquist or output filter, bust support, integrated A/D or D/A, Standards supported, and final Eval Board availability.

The primary providers shown below are Stanford Telecom (STEL), VLSI Technology (VES), and LSI Logic. Broadcom is also a significant supplier of relevant technology, however the NDA they require for access to technical information essentially amounts to a non-compete agreement.

	Bd/Chip	QPSK	QAM-16	Equalizer	Data Rate	Scmblr	Diff Enc	RS	Viterbi	Interlvr	Filter	Burst	A/D-D/A	Stds	Eval bd
Modulators															
STEL-1109	chip	Y	Y	na	20/40Mbit	Y	Y	Y	N	N	Y	Both	Y	M,D,802	Y
L64767	chip	Y	Y	na	7.8Mbaud	Y	Y	Y	N	Y	Y	cont	N	DTVB	?
Demodulators															
CATV															
STEL-2105	chip	Y	N	N	10Mbit	N	na	N	N	N	Y	Cont	Y	*	?
STEL-9236	Bd	Y	N	N	8Mbit	Y	Y	N	Y	N	Y	Cont	Y	*	na
VES1520	chip	N	Y	32-tap	8.7Mbaud	Y	Y	Y	N	Y	Y	Cont	N	DVB-C	Y
L64768	chip	QAM-4	Y	Y	7Mbaud	Y	Y	Y	N	Y	Y	cont	Y	DVB	
Satellite															
VES1877A	chip	Y	N	N	30Mbaud	Y	Y	Y	Y	Y	Y	Cont	Y	DVB-S	Y
L64724	chip	Y	N	N	45Mbaud	Y	Y	Y	Y	Y	Y	Cont	Y	DVB-S	Y
CDM8511	chip	Y	N	N	45Mbaud	Y	Y	Y	Y	Y	Y	cont	N	DVB-S	Y
Modems															
VES1900	chip														Q1-98
STEL2176m	chip	Y	Y	na	7Mbaud?	Y	Y	Y	N	N	Y	both	Y		Y
STEL2176d	chip	N	Y	Y	7Mbaud	Y	Y	Y	Y	Y	Y	cont	Y		Y

Table 1: Off-the-shelf Technology Survey

The off-the-shelf technology that Travertine is focusing on are the Stanford STEL-1109 burst transmitter, the LSI Logic L64767 CATV QAM modulator, and the L64768 QAM demodulator, and the Stanford STEL-2176 integrated modem. The VLSI Technology VES1520 CATV QAM demodulator was rejected primarily because control of the chip is through a serial IIC port only, which is too slow to allow rapid modification of equalizer and other filter parameters.

The STEL-2176 is promising as a single chip .35 micron technology demonstration/first product vehicle. However, because integrated subscriber modems are not designed to communicate between themselves, compromises are necessary to use this part. QAM-16 is the only modulation mode in common between the transmitter and the receiver, and the receiver cannot turn off the convolutional decoding (and perhaps interleaving as well), both of which are features the transmit side of the chip do not support. This difficulty can perhaps be worked around by:

a) convolutionally encoding the data in software before applying it to the transmitter
b) taking the data out of the receiver prior to the FEC logic block, and performing the requisite FEC decoding in software

Both options necessitate the presence of a fast local controller.

As of the time of this writing, Stanford Telecom is examining options for use of the 2176 as technology demonstration vehicle as part of a joint development proposal.

L64767 and L64768 Pros

The LSI logic CATV modulator and demodulator chips form a matched pair capable of supporting 7Mbaud. Some of the desirable features include:

a) A byte wide microprocessor interface port for programming – important for programming speed considerations, particularly for the adaptive equalizer coefficients.
b) An external uncorrectable error signal and internal error count registers
c) The interleaver depth is programmable
d) An FEC bypass mode in which the uncorrected error stream is presented externally. This feature is very useful for diagnostics and channel analysis
e) The LSI/DVB demodulator expects input centered at the "low second IF" frequency of from 7 to 7.8 Mhz . *At a 7Mhz baud rate this would put the lower edge of the HAN passband in the 3.5 – 4.2 Mhz range – exactly where we are looking to place the lower edge of the HAN passband. Hence it appears we can operate from the modulator/demodulator directly, without any additional mixing*
f) The Nyquist transmit filter coefficients (108 coefficients) are completely programmable.

L64767 and L64768 Cons

a) These chips send data in DVB/MPEG-II transport standard sized packets with a 187 byte payload with a prepended sync byte and an appended 16 byte RS block for a total packet length of 204 bytes. The overhead of dealing with this packet size remains to be understood.
b) Convolutional FEC (Viterbi) is not supported. This leaves the higher latency RS coding and/or interleaving.
c) There was no design intent, and thus no explicit provisions for rapid reacquisition of new transmitters in a dynamic multipath environment – which means no support for on-chip storage and rapid loading of pre-trained equalizer, filter, and timing recovery coefficients and parameters.
d) The support of higher constellation QAM, along with the necessary increase in hardware and AFE (8-bit A/D, etc.) complexity are not needed.
e) The microprocessor interface is not PCI compliant, but is more ISA like.

Much To Do

Areas of intensive focus for Travertine include:

1) Building an empirical database of premises wiring characteristics. Methodology, metrics, test equipment to be used, etc.
2) A MAC layer must be fleshed out that meets the long term goals of the HAN, and yet can be supported by the selected off-the-shelf technology. There is risk either that the existing silicon cannot provide adequate MAC support, or that the adjustments that must be made to the MAC layer to accommodate the existing silicon will compromise future developments/capabilities because of the backward compatibility requirements. Existing specifications such as the MCNS Data-over-Cable RF Interface Specification will help in understanding what protocol capabilities can be supported by existing silicon.
3) Network management issues
4) Custom HAN modem specifications, and the selection of a silicon partner to begin work.

References

[1] W.E. Stephens and T.C. Banwell, *"Line Code Selection for 155.52 Mb/s Data Transmission on Category 5 Cable Plant"*, IEEE Journal on Selected Areas in Communications, Vol.13, pp. 1670-1683, December 1995.