BRIDGING AND MANAGING QOS - CABLE AND HOME NETWORKS

Mariano Schain Itay Sherman
Texas Instruments – Cable Broadband Communications

Abstract

In order to enable end-to-end Quality of Service (QoS) for different services, sush as voice and video, a mapping of the QoS mechanisms of any two bridged networks must be supported and configured. A direct mapping approach is not scalable or feasible. In this paper we present a generic QoS scheme. We present in detail the QoS framework of DOCSIS1.1 and provide an overview of the QoS support in current and future versions of home networking protocols and architectures (e.g. 802.11, HPNA). For the frameworks presented, we explore the method to map their QoS mechanisms to the generic scheme, thus allowing scalable indirect QoS bridging.

INTRODUCTION

Cable operators are facing a new challenge in distributing their services into the home.

The DOCSIS (Data Over Cable System Interface Specification) standard has been widely adopted by the cable industry as the prevailing protocol for Internet over cable. The first version of the specification (DOCSIS 1.0) mechanisms was targeted at basic data transmission, providing a best-effort level of service. The new version (DOCSIS 1.1) is built on top of the previous 1.0 specification and enables operators to provide consistent and reliable digital services (such as Voice and Video) through the use of sophisticated QoS (Quality of Service) and Network Management.

However, with the emergence of multiple Home Networking technologies, operators find themselves using technologies over which they have significantly less control to distribute their services throughout the home. While these technologies can enhance the customers' experience, they also potentially put the reliability of cable services at risk. To further complicate things, many Home Networking technologies are being offered to consumers: wireless, phoneline, powerline and new wire technologies. In most of these categories, several technologies are competing with each other. Furthermore, the list of competing technologies is getting longer as new technologies emerge. It is expected that several Home Networking technologies will find their way into the home, resulting in a heterogeneous home network.

In such a heterogeneous environment, a data flow carrying service information may cross different home network segments on its end-to-end path. The service quality can't be guaranteed unless it is configured for each and every segment. Special consideration is given to the bridging points – the network entities that attach different home network segments together. These network bridges are responsible for forwarding the data flow across their network interfaces so that it seamlessly crosses the segment's boundary without affecting the quality of service.

To achieve the QoS-bridging goal, one may attempt to map the QoS mechanisms of any two protocols 'to be bridged.' This approach, however, is not always possible -

different network protocols provide different levels of QoS. This approach also is not always efficient because too many parameters must be considered in the mapping. Moreover, the scalability of such an approach is more than questionable because, for every new network protocol, a mapping to all existing protocols should be provided.

A generic QoS specification, as described in this paper, is both feasible and scalable as no direct mapping is required. Every network technology is mapped to the generic QoS specification, requiring a linear total number of mapping rather than square.

The paper is organized as follows: first, we present in detail the QoS framework of DOCSIS1.1. Next, an overview of the different home-networking technologies is given, including their strengths and weaknesses and for which applications they are most well-suited. The next section presents the concept of generic QoS specification and then an example mapping is given for DOCSIS and some of the other home networking architectures. We conclude with the evolution of the home networks toward a unified QoS architecture.

DOCSIS

Cable Modem (CM) and Cable Modem Termination System (CMTS) are the main entities in the DOCSIS network protocol. Several CMs- residing in the customer premises - are connected to a CMTS - residing at the cable operators' Head-End - through a Hybrid Fiber-Coax (HFC) network in a 'tree-like' structure - where the Head-End is the root and the CMs are at the leaves.

The DOCSIS downstream (DS) channel - carrying information from the CMTS to the CMs- uses a typical TV channel (6MHz wide in USA, 8MHz wide in Europe) to carry 'Ethernet-like' packets over a continuous digital MPEG stream. As the DS channel is shared, the channel bandwidth (about 40 Mbps in USA) is distributed among all active CMs. Note that all the DS packets are received by all connected CMs. The Ethernet address is used by the CMs to filter out the packets they need.

The DOCSIS upstream (US) channel carries 'Ethernet-like' packets and uses the lower frequencies (below the range allocated to TV channels), a range that is prone to occasional interference. To cope with such situations, the US channel configuration (bandwidth, rate, error correction, and other transmission parameters) is diverse and is dynamically controlled by the CMTS. The transmission scheme is decided on a burst-by-burst basis. Again, the US channel is shared among several CMs, hence a multiple access mechanism is implemented and the channel bandwidth (up to 10Mbps in the current specification) is managed and allocated to active CMs by the CMTS. Note that US packets are received by the CMTS only. Packets from CM to CM always pass through the CMTS.

A DOCSIS domain may include several downstream and upstream channels paired accordingly to achieve the required network balance. CMs may be instructed by the CMTS to move from channel to channel as a load balancing implementation or as a means to overcome channel quality problems. A CM acts as a transparent bridge, it forwards packets that are received from the CMTS toward its local network interface (Ethernet, USB, etc') and vice-versa. Packets that are destined to the CM (such as packets to the SNMP, DHCP or other IP based agents residing in the CM) are consumed by the CM and not forwarded.

As CMs reside in varying distances from the CMTS, The Time Division Multiple

Access (TDMA) scheme implemented in the US requires subtle synchronization mechanisms. All CMs align to the CMTS clock (this clock is distributed through dedicated DS control messages). A ‘Ranging’ mechanism (in which the CMTS instructs the CM on the time shift and power level to use in US transmissions) is constantly active for every connected CM. This mechanism ensures that all CMs transmissions are aligned to a time base controlled by the CMTS, and that all signals are received at the CMTS at approximately the same level, ensuring the ability of the CMTS to identify collisions.

The US channel is divided into time slots. A transmission interval is a group of continuous time slots. The CMTS allocates transmission intervals for different needs (transmission requests, packet transmission, ranging messages, etc.), and transmits the allocation to the CMs using a dedicated control message in the DS. Some of the allocated intervals are multicast (generally, transmission request intervals are not allocated to a specific CM) and may result in a contention, and some of the intervals are unicast (such as a packet transmission interval). The US interval allocation messages (‘MAP’ messages) are transmitted in the DS by the CMTS in a timely manner. To illustrate the US transmission mechanism, consider a CM wishing to transmit a packet in the US. The CM will analyze the MAP messages in the DS until a multicast ‘request’ interval is allocated by the CMTS. The CM will transmit its request in the specified interval (the request contains the required transmission length) and will wait until a MAP message containing the requested allocation (a ‘grant’) is received. Once the grant is received, the CM transmits the packet in the allocated interval. It might be the case that more than one CM transmitted a request in the same interval. In that case, a back-off algorithm is implemented to solve the contention. To overcome the possible contention of subsequent requests, a CM may transmit a request embedded (‘piggybacked’) in a packet transmission.

While the best-effort scheme that was provided by DOCSIS 1.0 was sufficient for basic Internet access, it fell short of the needs of more sophisticated services that are unable to operate in the absence of guaranteed QoS. The QoS framework of DOCSIS1.1 (as detailed below) was targeted at exactly those types of services.

Four main service categories are supported by DOCSIS1.1: Unsolicited Grant Service (UGS), Real-Time Polling Service (rtPS), Non-Real-Time Polling Service (nrtPS), and Best Effort (BE) service.

In a UGS flow, the CM is assured to receive from the CMTS fixed size grants at periodic intervals without the need to explicitly send requests. In addition to the grant size and the period, the tolerated grant jitter is also negotiated at service setup. The main advantage in using a UGS is the reduced latency achieved by eliminating the need to go through the request-grant cycle for every packet. However, using a UGS is inefficient for applications that don’t require a constant data rate over time. A flavor of UGS - Unsolicited Grant Service with Activity Detection (UGS-AD) – is targeted at those exact applications (e.g. Voice with silence detection). In a UGS-AD, once the CMTS detects flow inactivity (through non usage of grants by the CM), it starts sending unicast request opportunities (also called ‘Polls’) at periodic time intervals. The CM can use the unicast requests opportunities to send requests (once the flow is to resume) avoiding the latency incurred by contention at multicast request intervals.

In a rtPS flow, the CM is assured to receive from the CMTS unicast request opportunities at periodic intervals. If the CM does not use the request opportunities, the CMTS allocates the reserved bandwidth to other flows, overcoming the inefficiency of UGS. In a nrtPS flow, the bandwidth is not guaranteed to the flow. The CM, however, is allowed to use multicast request opportunities for the flow as well. The last QoS category, Best Effort, defines the minimum traffic rate (which the CMTS must reserve) and the maximum allowed rate as the main service parameters.

Multiple data flows (each flow corresponding to a service and identified by a Service ID (SID) may concurrently exist in a CM. A transmission request in the US and the corresponding grant includes the SID as the flow identifier. The CM and the CMTS negotiate the QoS for each flow upon allocation and dynamically as the service requirement changes (in dedicated procedures). The QoS is then achieved by the implementation of sophisticated scheduling mechanisms in the CMTS. A classification function is applied to every packet. The flow in which a packet is transmitted is based on the content of the Ethernet and IP header fields (allowing every application to receive a different service flow). The classification function may also indicate the suppression of the packet header (a mechanism that is useful for packets with semi-constant headers and short content, such as voice packets).

802.11

IEEE 802.11 are a set of standards for wireless networking. The original 802.11 specification, defined a network that works in the 2.4Ghz frequency band and is capable of transmission up to 2Mbs. The 802.11b standard that is becoming a dominant force in the wireless home networking arena extends this capability to 11Mbs thus providing the ability to transfer high quality audio and video content over the network. 802.11b PC cards, PCI cards and external base stations that are sold today enable multiple PCs to connect with the home or office as well as other IAs such as PDAs. The 802.11a standard provides even higher data rates (>50Mbs) using the 5.7Ghz band. The 802.11g specification will provide a direct extension to the 802.11b using the same 2.4Ghz band and providing data rates of 22Mbs.

An 802.11 network is usually composed of multiple stations and a single access point that coordinates the network activity. In some cases an Adhoc network can be established with stations only.

The original 802.11 specification used a basic contention-based MAC and as such does not provide real mechanisms for QoS (similar to plain old Ethernet). The ongoing work on the 802.11e specification is intended to provide the necessary extensions to the 802.11 MAC in order to provide real QoS. It is important to note that the original 802.11 MAC spec took in to consideration the possibility of MAC extensions. Therefore, old 802.11 devices with the original MAC will not degrade the QoS capabilities of newer 802.11e devices. Hpwever, this is not the case for Ethernet as well as Home Networking Phone Alliance (HPNA).

The 802.11e specification includes multiple levels of QoS. The basic level provides for a simple prioritization mechanism similar to the basic HPNA 2.0 mechanism and 802.3p. The higher levels of 802.11e are targeted to provide real QoS with guaranteed bandwidth and delay on a per-stream basis. These capabilities are similar and, in some cases, exceed the ones that exist in the DOCSIS 1.1 specification.

The 802.11e specification includes multiple bandwidth management schemes. The lower levels employ a distributed approach where each station on the network is making its own QoS decisions similar to Ethernet. The highest level (3) uses a centralized approach where the accesses point controls the assignment of bandwidth to the client units similar to a DOCSIS CMTS and DOCSIS CMs. All these schemes can coexist on the same physical channel. The coexistence of different level QoS devices is similar to the co-existence capability of DOCSIS 1.0 and DOCSIS 1.1 modems on the DOCSIS network.

The 802.11e draft defines a set of MAC sub-layer QoS parameters more extensive than may be needed, or may be available, for any particular instance of QoS traffic. These parameters, collectively called a traffic specification and applied to a traffic category (TC), are Traffic Type, Ack Policy, Delivery Priority, Retry Interval, Polling Interval, Transmit Interval, Nominal MAC Service Data Unit (MSDU) Size, Minimum Data Rate, Mean Data Rate, Maximum Burst Size, Delay Bound, and Jitter Bound.

HPNA

The network defined by the HPNA enables devices to transmit packets over residential phone lines. Two specifications have been developed: HPNA 1.0, which offers a data rate of 1 Mbps using an Ethernet-like Carrier Sense Multiple Access (CSMA) protocol without any QoS mechanisms, and HPNA 2.0, which offers a data rate of 10 Mbps using a prioritized CSMA protocol. In the discussion that follows, we refer to the later HPNA 2.0 protocol.

The management of the PHY and MAC layers is distributed. All the devices cooperate to enable minimal overhead of network resources usage for network management. All decisions, such as transmission schemes and priority, are implemented locally in every device. For example, every device decides, on which transmission scheme to use based on the destination device(s) characteristics.

The network supports eight priority levels from zero (lowest) to seven (highest). A time-division scheme is used to enable prioritized access. For every priority level, a designated time slot is defined (the time slot for priority seven being first, followed by the time slots for the other priorities in descending order). Devices are allowed to start the transmission of a packet only in (or after) the time slot that corresponds to the packets' priority (as decided by the device). Since collisions can still occur - usually between packets of the same priority - a contention resolution algorithm (CRA) - based on a random back-off mechanism - is implemented. New packets for transmission do not preempt the CRA unless they are of higher priority.

The above scheme carries some inherent characteristics that impact the usability of the HPNA network for services requiring guaranteed QoS. Ignoring collisions, the worst-case latency of a packet to be transmitted is in the range of the maximum transmission time of a packet – about 4.2 milliseconds. If collisions occur, however, the latency can be only statistically bounded (every contention resolution 'round' takes about 200microseconds). As the stations are free to choose the transmission priority of most of the packets (although a mapping of IEEE 802.1D priority levels to HPNA priorities is provided), the resulting network latency can't be guaranteed (higher priority 'preempting' packets may affect the CRA completion time).

The above limitations call for higher-level network synchronization - a management entity at a higher layer that controls both the priority and timing of packet transmission at the MAC and PHY layers. Such a scheme will enable the provision of quality demanding services such as voice and video over phone lines.

IEEE1394

The IEEE1394a standard defines a wired networking targeted to connect very-high-speed multimedia devices over short distances. The IEEE1394a supports connection speeds of up to 1.6Gbps with lower levels starting at 200Mbps. IEEE1394 uses special types of cables to carry the information. Multiple IEEE1394 devices may be connected to each other in a daisy chain formation. IEEE1394 interfaces are part of every digital camcorders and are also widely deployed in multimedia PCs. It is also used to connect high-end video display devices (HDTV) and editing equipment. The IEEE1394 was also chosen as the interface between the next generation of OpenCable™ set top boxes and HDTV displays. Efforts are being made to extend the capabilities of IEEE1394 and to provide for wireless variations. Other IEEE1394 related specifications define the way MPEG-2 TS (the standard used for Digital broadcast over cable and satellite) are carried over IEEE1394 networks, delivery of IP traffic over IEEE1394 and control language over video equipment (start, play, stop, forward, back, titles etc.).

The IEEE1394 MAC layer includes two types of channels: asynchronous and isochronous.

The asynchronous channel is a best-effort data channel that has no QoS mechanism associated with it.

The isochronous channel is a reservation-based, assured-bandwidth channel.

A designated isochronous resource manager in one of the devices connected to the IEEE1394 network manages the Isochronous channel. The resource manager accepts reservation requests from all devices on the network and assigns specific bandwidth to each one of them. The reservation should reflect the peak rate of transmission from the source device. Bandwidth that is not assigned to any Isochronous reservations or is not utilized by an existing reservation is used to carry best-effort asynchronous transmissions.

It is important to note that the capacity of the DOCSIS channel is negligible compared to the capacity of 1394, which is typically 200-400Mbps and the garnered capacity to a subscriber using DOCSIS channel is a fraction of the overall 40Mbps D/S channel). Also the jitter and delay provided by IEEE1394 on the isochronous channel is extremely low due to the high bit rates and fixed reservation mechanism.

BlueTooth

The Bluetooth (BT) specifications define a short-range wireless network working on the 2.4Ghz band. The maximum data rate supported on a BT system is 768Kbs. The BT specification defines two connection methods: The point-to-point connection where two devices are creating a private connection, and the point-to-multi-point connection where a real network of up to eight active devices is created. The point to multi-point connection is based on an assignment of a master device to each such network (named a piconet) and multiple slave devices. A slave device may belong to multiple piconets. BT is rapidly deploying today in devices such as cell phone PDA etc..

The definition of mapping in this section relates to the point-to-multi-point connection defined in BT.

BT defines two types of channels for transmission - synchronous and asynchronous. Synchronous channels (SCO) are symmetric and provide a 64kb/s bi-directional connection between the Master and a specific slave. Transmission and receiving slots are sent periodically with a fixed interval between them, each slot is 625usec in length. Up to three such slots can be accommodated by a single master or slave. The SCO slots are setup by the master using the BT LM protocol. There are multiple types of SCO slots, most of them are targeted to voice distribution, but one of them can carry both voice and data simultaneously.

Asynchronous packets (ACN) are sent on the slots left after SCO assignment. The slaves send information only after they receive information from the master. There are multiple types of ACN slots that differ by their payload size and FEC protection. For the ACN channels, the BT specification defines the L2CAP layer that enables segmentation and re-assembly of packets as well as QoS services and connection establishment.

The SCO channels of BT are really targeted at voice streams for which they provide QoS. The different types of slot VH1..3 provide different packetization periods and delays for the voice transmission. It is important to note that these slots are highly tailored for simple PCM-based transmission of voice packets. Usage of these slots for other applications or voice codec types is not trivial. The usage of SCO channels is also limited since the BT L2CAP conversion layer is not defined for this channel and thus the direct translation to IP is not obvious. It is most likely that SCO channels will be used for voice delivery only.

Best-effort data as well as reservations not using SCO will use the ACN slots. The L2CAP layer of BT runs on top of ACN channels and includes a QoS reservation protocol that uses values, such as bucket rate and jitter, to define the QoS of a designated stream. The L2CAP QoS is not mandatory and so not all BT implementations support it.

From the QoS perspective, BT SCO channels provide for excellent QoS but are intended for a very narrow set of voice applications. The L2CAP QoS mechanism can potentially provide for good QoS but it is not mandatory as part of the spec. The overall bandwidth capabilities of BT limit it's ability to support high-end applications such as mid/high quality video streaming, as such it can extend DOCSIS QoS for a limited set of applications.

HomeRF

The HomeRF specification is another standard for wireless networking. HomeRF devices are capable of transmitting up to rates of 1.6Mbps. Future HomeRF specifications are targeted to reach the 10Mbps range. HomeRF devices marked the first phase of commercial wireless home-network devices. HomeRF PC Cards, PCI cards and external modems that are being sold today enable the connection of multiple PCs and peripherals with in the home.

The HomeRF MAC layer architecture is a combination of an asynchronous, Ethernet-like access mechanism and an isochronous, circuit-switched TDMA access mechanism.

The isochronous channels are used primarily for up to eight active 32 kbps ADPC voice connections using DECT signaling in the upper layers. As in BT SCO channels, these channels are providing full QoS with guaranteed delivery and bounded jitter and delay, but only for very narrow range of voice applications.

The asynchronous data service is comprised of both a best-effort service and a prioritized asynchronous service. The management of the PHY and MAC layers is distributed. All the devices cooperate to enable minimal overhead of network resources usage for network management. All decisions, such as transmission schemes and priority, are implemented locally in each device

HomeRF devices support eight prioritization levels for asynchronous data as well as a lower level of best-effort data.

Specific streams can be assigned higher priority over other streams. It is also possible to define a nominal time reservation and maximal time reservation for a stream.

Like 802.1p, the HomeRF asynchronous channel can only provide for prioritization but with no guarantee of delay or bandwidth.

Its limitations are similar to the ones identified for HPNA 2.0.

The current spec of HomeRF is also limited in the application it can carry due to its bandwidth. It would not be suitable to carry mid/high quality video streams.

HomePlug

Although power-line, as a communications medium, is more complex than wireless and phone-line, the distribution of outlets within the home makes it a viable platform for home networking. The HomePlug Powerline Alliance was formed to create an industry standard for high-speed home networking over power lines. Similar to the situation with phone-lines, the management of the PHY and MAC layers is distributed - all decisions are implemented locally in each device depending on a dynamic configuration that is suited to the specific pair of interconnected outlets. The adaptation to specific channel characteristics is essential to the efficient use of the media since the exact channel behavior is unique and generally unpredictable. An encryption scheme in the MAC layer copes with the inability of the distribution transformer to block the propagation of power-line signals from one home it powers to another.

A variant of the Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA) shared access protocol is used (with a random back-off algorithm for contention resolution). QoS is supported through the additional definitions and mechanisms for priority classes and latency control: stations use a priority resolution interval to signal the requested transmission priority, ensuring the early transmission of the higher priority frames. Latency is reduced through the use of segmentation and reassembly (a mechanism is provided to allow uninterrupted transmission of multiple segments in the absence of higher priority frames). By enforcing the discard of stale packets, a bound on the latency is achieved. This is particularly important for applications demanding latency control (e.g. VOIP).

HomePlug shares many of the phone-line protocol limitations, primarily the inability to prevent a non-cooperating - though compliant- device with ‘private’ priority semantics. To truly support latency and jitter sensitive applications, higher layer coordination is a desired scheme in this case as well

QoS Based on flow-specs requirements

We are interested in the bridging of QoS from the DOCSIS technology to other home-networking technologies. This raises a question regarding what is actualy being bridged. As can be seen from the previous sections of this paper, there are multiple ways to achieve QoS over different networks technologies, some use CBR mechanisms (Constant Bit Rate) where a set of two numbers (Packet rate, Packet Size) define the QoS, while others use prioritization mechanism where a single, different value (Priority) defines the QoS. Other mechanisms use even more complex representations for the QoS with multiple values. Trying to directly match these QoS mechanism and values is in most cases impossible.

A different approach to QoS mapping would be to look at the actual streams of data for which we are attempting to provide QoS. We can use some sort of general QoS requirement definition for these streams, and then try and map these general requirements to any one of the different technologies on the stream’s way from source to destination. The mapping of two technologies is indirectly performed by mapping each to these general stream requirements.

One general model that had been used in multiple standards is the flow model, which imitates a fluid-flow model. In this model, the data flows are represented as buckets of fluid (the data is the fluid) that are purred into pipes of variable size. The flow is defined by the size of the bucket, the number of buckets per second, the peak rate for which a fluid can purr from the bucket and the pipe size. Additional values such as the over all delay can be added to provide for more information on the requirements of the stream. Although it seems very crude, the flow model has been proven to be a viable tool to defineing the real QoS requirements of data streams. These general flow requirements are titled "flow spec."

The benefits of this indirect mapping are not obvious at first glance. In order to better explain them let's consider the following example:

Technology A uses a QoS mechanism that is based on the assignment of a constant bit rate (CBR) channel for each stream. The channel is defined by the packet size transmitted and the time-gap between each of these packets, which is completely fixed.

Technology B uses a prioritization mechanism with limitation on packet sizes. Streams with higher priority are given access to the channel before streams with lower priority. The maximum length of a packet to be transmitted for each stream is also limited.

We wish to deliver the following stream that is defined by the flow model as follows:

Bucket Rate = 100 per second

Bucket Size = 100 bytes

Peak Rate = 10Kbytes/per second

Pipe bandwidth = 10Kbytes/per second

Max Delay = 30msec

Technology A may map these stream requirements to a CBR channel of packet size 200 bytes and repetition of 50 times per second. This definition satisfies the requirements of the stream.

Technology B would have to consider the other existing streams on the network and consider their relative priority and packet sizes. Let's assume that the network data rate is 100K byte per second, the number of priority levels is four, and three streams of level three already exist with a packet size of 2000b ytes each. The new stream will be mapped to level four with a packet size of 100 bytes. Mapping it to lower levels will cause a violation of the Max Delay defined by the stream.

We were able to map the stream to both technologies and, by that, create an indirect mapping between two completely different technologies. This mapping will enable us to deliver the specific stream over a network that is comprised of segments from both of these technologies while maintaining the stream requirements. However, we cannot assume that the direct mapping exists between the two. A "technology A" CBR channel of 200bytes and repetition of 50hz is not equivalent to "technology B" level four. The mapping exists only in the context of the specific streams that are active in the system on a specific time.

Mapping to Flow Specs

This section provides a short explanation of the mapping between flowspec parameters and some of the actual technologies previously described.

DOCSIS

The DOCSIS QoS is using multiple types of services such as CBR, Real Time Pole (RTP) etc.. The information provided by the flowspec can be used by the DOCSIS CMTS to decide what type of service a specific stream will receive. Streams that have their peak rate equal or close to their average rate will receive a CBR service and the parameters of this CBR will be set according to the peak rate of the flow spec, other streams that have a more bursty nature may use an RTP service. A more detailed mapping between DOCSIS parameters and flow specs can be found in CableLabs™ PacketCable™ QoS specification. It is important to note that the mapping information provided for DOCSIS is just a guideline for implementation. Every CMTS vendor can decide on different mapping algorithms. This does not create any interoperability or performance problems beacause DOCSIS uses a fully centralized bandwidth and QoS management.

802.11

The QoS mechanisms defined in 802.11e use parameters that are very similar to the ones of the flow specs. Therefore, a direct mapping between the two can be made. The physical mapping of these requirements to the channel behavior depends on which level of 802.11e QoS is used. Some of the levels use a distributed approach where each station is making it’s own decisions on the scheduling of packet for transmissions, while others use a centralized approach where a master station makes the decision for all other stations. In general the centralized approach can provide full compliance to the flow-spec requirements, while the distributed approach is usually statistical, and can provide a very high chance of compliance to the requirements, but not complete certainty.

In most cases, an 802.11b/e network can meet flow-specs requirements that can be met by the DOCSIS channel.

BT

There are multiple option to providing QoS over BT. In some cases, the flow-spec requirements can be satisfied by using an SCO channel (given the stream bucket size and rate fit the size of the SCO slots), other flow specs may be mapped to the L2CAP QoS parameters. These parameters are very similar to the flow-spec parameters and a direct mapping is possible. The scheduler in the master station performs the physical mapping of these requirements to the channel behavior. It is important to note that, due to the limited bandwidth of the BT network, it is likely that some streams flow specs that the DOCSIS network can satisfy would not be met by the BT network (i.e. a flow spec of data rate higher then 768kbit/sec). For BT devices that do not include the L2CAP QoS extensions, it is generally impossible to assure the fulfillment of the flow-spec requirements. If the rates of data in the flow spec as well as the delay requirements are significantly lower then the over all BT channel capacity, then there is a good chance that the flow-spec requirements will be met. But for rates that are closer to the channel capacity or strict delay requirements, the chances are that the non QoS BT network would not be able to satisfy these requirements.

HPNA

The current HPNA specification (2.0) uses a priority based QoS mechanism. The exact usage of the priority mechanism is not defined in the HPNA specification and is left open. Flow specs with higher bandwidth and delay requirements may be mapped to higher priority while best-effort or low-capacity/ high-delay flow specs will be mapped to lower priorities. This type of mapping can guarantee the requirements of the flow specs when we are dealing with a low number of streams. When the number of streams increases, we have a situation where multiple streams use the same priority level and it is impossible to guarantee the delay that packets from each of these streams may encounter. Using statistical models we can in most cases assure to a very high probability that the flowspec requirements will be met even in cases where multiple streams are using the same priority.

Adding a control protocol that will negotiate scheduling information between the different devices on the same HPNA network can make an improvement to the QoS provided by the basic prioritization mechanism. Such protocols are currently under investigation in different standardization groups.

It is important to note that all the described mechanism for QoS over HPNA require that all devices on the network will obey a defined set of rules for the usage of the priority mechanism. These rules are not part of the HPNA specification itself and since the scheduling of packets transmission is distributed (each HPNA device makes its own decision on when to transmit, there is no central point of coordination), it is hard to assure the behavior of a network that is using equipment from multiple vendors. The most obvious example of this problem is the connection of an HPNA adaptor to a PC where the priority level is set by the PC owner to the maximum (7), and large packets are sent at a high rate over this interface. Even if all other devices on this HPNA network are attempting to cleverly schedule streams to priorities, they will all be blocked by the transmission originating from the PC.

HomeRF

The HomeRF QoS mechanism is a combination of CBR channels and a priority based QoS channel. Similarly to BT, flow specs that fit into the HomeRF CBR channels may use it for delivery of streams (very narrow usage for voice streams). Other flowspecs will need to be mapped to the different priorities. The problems associated with mapping to a priority based QoS are explained in the HPNA section. These problems are significantly more severe when dealing with HomeRF as the capacity of the HomeRF channel is much lower then the one offered by HPNA 2.0. As in the case of BT, some of the flowspecs delivered by DOCSIS would not be deliverable over HomeRF because of its limited bandwidth.

IEEE1394

The IEEE1394 QoS is based on assignment of CBR channels for each stream. The mapping from flowspec requirements to CBR can be based on the R (pipe bandwidth) parameter of the flowspec. This is a very crude and simple mapping that does not take into account the other flow-spec parameters. But, for all practical applications, the other requirements of the flow spec will be met since the IEEE1394 channel overall bandwidth is large and the delays are short. This is specifically true when talking about streams that need to be carried over the DOCSIS channel that offer smaller bandwidth and longer delays. In general, streams that are carried over DOCSIS networks should not have problems being carried over IEEE1394 networks.

Conclusion: The evolution toward a unified framework

The current home networks are mainly used to share Internet connection and a printer. However, they serve as the baseline for the future home networks. As more applications, appliances and services become available, the underlying components of the future home network will be aggregated, resulting in a heterogeneous home network.

In today's home network, the home appliances, such as TVs, phones, PCs and, PS2, are generally owned and controlled by the home resident. Some of those end points are 'general purpose' in nature, meaning they are not dedicated to a specific service. Rather they may be used for different services upon need. Some others are dedicated to a service and, as such, are more prone to end up being owned and controlled by the service provider. The home network itself, connecting the end devices to the service sources and making the services possible, is a clear candidate to be handed-over (in terms of management and control, rather than ownership) to the responsibility of an entity other than the home resident. The diversity of configurations will make it impossible for the home resident to support and maintain it by itself.

A unified framework is thus required to enable external control of the home network and a way to analyze the network topology, its components and their capabilities. For every network component, a method is required to control and configure it to other components in order to enable the service provision and assure the quality of service end-to-end. The components of such a framework are being defined today in several standardization efforts (OSGI, UPNP, CableHome et al.). Some of those standards complement each other, and some compete. Only a set of standards that solves all the aspects of controlling the home network (QoS, security, management and configuration, etc.) will prevail.

The CableHome™ specifications defined by CableLabs, aim at precisely that. The concept of generic flow specs lies at the heart of the CableHome model, so as to incorporate IP- based mechanisms such as SNMP, SBM, RSVP and others. However, to prevail as the 'end of the road' framework (and this applies to other candidate frameworks as well) special attention must be given to supporting an evolutionary process. Both existing components and those to be aggregated to home networks while the frameworks converge will have to be supported.

Allowing an evolution rather than imposing a revolution is the main challenge of the future home networking framework.

.