

The Multimedia Over Coax Alliance

This paper presents a historical, market, and technical overview of the MoCA home networking technology.

By ANTON MONK, RONALD LEE, AND YOAV HEBRON

ABSTRACT | The Multimedia over Coax Alliance (MoCA) publishes a home networking specification that is the deployed standard for the majority of service providers in North America. It is also being adopted in other parts of the world. MoCA enables highly robust, secure, and low-latency communication at net throughputs of over 400 Mb/s on the home coaxial cable. Its primary use is for the distribution of premium content video, including applications such as multiroom DVR. This paper presents a historical, market, and technical overview of MoCA.

KEYWORDS | Access protocols; Communication standards; multimedia communication; OFDM; home networking; Multimedia over Coax Alliance (MoCA)

I. INTRODUCTION

Home entertainment networking technology standardized by the Multimedia over Coax Alliance (MoCA) is currently being deployed in tens of millions of homes throughout the United States by cable, satellite, and telco service providers such as Comcast, Cox Communications, Time Warner Cable, DIRECTV, DISH Network, and Verizon. Internationally, MoCA is being deployed in Europe by Liberty Global, in Latin America by DIRECTV PanAmericana, and interest is increasing rapidly in China. The technology enables a high-speed, robust and quality-of-service (QoS)-based network on the in-home coaxial cable, allowing distribution of multiple high-definition (HD) video streams and broadband data to every room in the house. Key applications of a MoCA entertainment network include multiroom DVR (mDVR), PC-to-TV connectivity, broad-

band data service extension from an external gateway into the home, gaming, over-the-top (OTT) services such as Netflix and Apple TV, and Wi-Fi coverage extension leveraging MoCA as a networking “backbone” in the home.

Current deployments by service providers are based on the MoCA 1.1 standard [1], ratified in 2007, which enables net throughput of 175 Mb/s throughout the home and a guaranteed bandwidth approach to QoS. While this performance meets all the current needs of service providers, the new MoCA 2.0 specification [2] sets the stage for future services and applications in the home by extending performance beyond 400 Mb/s, adding many additional features needed to support those services, and providing critical full backwards interoperability to the many millions of deployed MoCA devices.

This paper provides an in-depth look at the history of in-home multimedia distribution that led to the creation of the MoCA standard. We also examine the market and applications that have driven the mass adoption of MoCA. In the main portion of this paper, we go into the technical details behind the various MoCA standards and how service providers are deploying the technology. Finally, we look at some other standards that, together with MoCA, encompass some of the ecosystem of technologies being deployed in the modern in-home digital entertainment world.

II. HISTORY

While there are many drivers of a home multimedia network, the original primary use case was mDVR. As this feature became widely deployed, the presence of the MoCA home network allowed operators to think about a gateway-client architecture which allows for the termination of the access network, broadcast video and services authentication at the gateway and conversion to an all-IP distribution inside the house to client devices (including consumer electronics devices such as tablets and televisions).

Manuscript received March 7, 2013; accepted May 22, 2013. Date of publication July 12, 2013; date of current version October 16, 2013.
The authors are with Entropic Communications, Inc., San Diego, CA 92121 USA (e-mail: Anton.Monk@entropic.com).
Digital Object Identifier: 10.1109/JPROC.2013.2266299

A. The DVR

In 1997, two companies—TiVO and Replay TV—revolutionized the industry by developing the first digital video recorders. Initially, uptake for the home DVR was slow, but as service providers started to adopt the technology in 2000, demand began to grow. Once digital recorders became available as part of their cable or satellite package, and users found how easy and convenient it was to record their favorite content, the demand for DVRs surged. Consumers began to appreciate that they could have access to all their content when it was most convenient for them to watch and that they had total control to pause, fast forward, rewind, and preset repeated recordings of all their favorite programs without the hassle of tape or DVDs.

Today more than 52% of all pay TV households in the United States have at least one DVR¹ and the traditional ways of viewing television is becoming obsolete as viewers are able to watch their favorite programming on their schedule. Demand has shifted from wanting DVR service, to wanting multiple DVRs in the home (see Section II-B). In the United States, 43% of households with DVR service have more than one DVR and that number is growing rapidly. This has created a new challenge for service operators as the cost of deploying multiple DVR units is very high.

B. Multiroom DVR

One of the biggest DVR challenges for consumers has been where to record and store content. When a consumer wants to record a show and have it available for viewing in any room, he has to program the DVR in each viewing area to record the same program so that it can be watched on the TV of choice. This scenario is time consuming and can be frustrating, especially when the consumer forgets to record on each DVR. In addition, the service provider is stuck with providing (and subsidizing) multiple DVRs per home.

The solution is a single, centralized content repository in the home: content stored on one large hard drive that is available for streaming to any device in the home. Devices are networked inside the home and consumers have total control over how and where they can access their in-home content and are able to stream that content to any device (assuming all content rights are met).

The basic requirement for achieving this is a very high-speed, robust IP-based network in the home—essentially a home entertainment network. Such a network needs to be able to handle real-time delivery of multiple video streams, including standard definition (SD) and HD and associated trick modes (fast forward/rewind), which can require significantly more bandwidth than the original stream. In addition, such a network needs to have an

¹Leightman Research Group.

inherently low packet error rate (PER) since there is no time to retransmit video packets. Another key feature is low latency: consumers have come to expect almost instantaneous response from their actions. Latency is also a critical requirement for gaming and for in-home content protection, as will be described in Section IV-E3.

Other key benefits of this architecture include lower overall power dissipation (since the IP client devices have no hard drive) and reduced overall system cost.

C. Early Home Networks (HPNA, Wi-Fi, HomePlug)

Moving video content around the home over a high-speed, robust, and reliable IP network has been a service provider industry goal for well over ten years. Many technologies have been applied to the problem and most have failed. The reason is that a key cost-driven system requirement for service providers is for no new wires (such as Category 5 Ethernet cables) to be added to the home, and for minimal remediation of existing cabling and infrastructure. Wireless technologies have not proven to be sufficiently robust or have high enough throughput and coverage to meet the stringent requirements of HD digital video. Of the existing standards for in-home cabling, both HomePlug/IEEE 1901 (the primary powerline networking standard) and the Home Phoneline Networking Alliance (HPNA) have also proven unable to meet these requirements, which were originally typically based around a minimum of 100 Mb/s of net throughput over 95% outlet coverage, with simple remediation for the remaining 5%.² While raw physical layer data rates continue to rise for these technologies as their standards are improved, the overall coverage in a home has remained too low to be useful for ubiquitous deployment of premium home entertainment networking. The remaining wire in the house, with the potential to carry vast amounts of video content, unsurprisingly perhaps, is coaxial cable, which also reaches most televisions in the home. Every major service provider in the United States has settled on coaxial cable as the primary way to move their traffic around the home.

D. Multimedia Over Coax Alliance

MoCA is an industry standards alliance formed in 2004 to specify and promote a coax networking technology developed originally by Entropic Communications, Inc., San Diego, CA, USA. The alliance also certifies products by specifying a Certification Test Plan and conducting interoperability tests of all MoCA products at an authorized independent testing facility. The Promoter members (represented by the board of directors) of the alliance include service provider representation from all major market segments (cable, satellite, and telco). This is

²More recently, these coverage numbers have increased to over 300 Mb/s of net throughput at 95% coverage, necessitating the move to the second generation MoCA standard.

because MoCA uniquely supports all these segments due to its frequency spectrum flexibility, as will be explained in Section III-A.

III. MARKET AND APPLICATIONS

A. Service Providers

The Pay TV market in the United States encompasses about 100 million subscribers. Most service providers have spent decades and billions of dollars developing the infrastructure to support high-quality HD video and triple-play services that, in many cases, justify an average revenue per user (ARPU) of more than \$100. These service providers have large call center operations to handle subscriber calls for any issue related to Internet, telephone, or video services. This is very different from a typical consumer electronics provider. As such, service providers are uniquely sensitive to the quality and consistency of experience they deliver and are unwilling to jeopardize the subscriber base with an in-home network choice that lowers the quality of premium and HD content. This is one of the reasons why it took so many years for service providers to deploy mDVR services—wireless, powerline, and phoneline solutions had unacceptably poor coverage. With the commercialization of MoCA's coax solution, delivering robust, high-quality and ubiquitous in-home coverage, service providers were confidently able to

extend their network and IP infrastructure into and throughout the home.

The three segments of service providers (cable, satellite, and telco) use different MoCA architectures and/or frequency bands. A typical cable operator home is shown in Fig. 1. Broadband data and video is delivered over the coaxial cable plant external to the house in frequency bands below 1 GHz. The external plant is typically connected to all coax outlets in the home. In the example below, a DVR receives and may store content and then distributes that content over MoCA to other set-top boxes (STBs) in the home. A cable modem terminates the broadband data signal and may distribute that Ethernet data over MoCA to any Ethernet-connected devices in the home that have an Ethernet-to-MoCA adapter connected to the coax. In addition, the example shows a second Wi-Fi Access Point connected to the MoCA network to allow for greater in-home coverage.

Cable operators currently utilize a single MoCA channel above the video spectrum (which can extend to around 1 GHz). This channel uses the “D” band of MoCA, which extends upwards from 1125 MHz (see Section IV-E6b).

Satellite operators use a similar architecture but, because satellite signals are transported from the dish into the home at L-band (above 950 MHz) and these must coexist with MoCA, they use MoCA in a lower frequency band—either E-band (DirecTV) or F-band (DISH Networks). See Section IV-E6b for details. In addition, since




Fig. 1. Typical cable operator architecture.

satellite homes have a separate broadband data connection (Cable or Digital Subscriber Line) this connection is bridged onto the MoCA network through an Ethernet to MoCA adapter.

Verizon FiOS uses a somewhat unique architecture. For the in-home network, they have the same basic architecture as the cable operators. However, to extend the fiber access network from the optical network terminal (ONT) to the broadband home router (BHR) in the home, they utilize a second MoCA band in a point-to-point mode. This band is the C4-band (see Section IV-E6b). The ONT serves a similar purpose to the cable or DSL modem in terminating the broadband data connection and bridging the Ethernet traffic onto the MoCA network.

B. The Consumer

With MoCA networks now deployed in millions of U.S. households, it makes sense that consumers should be able to attach additional consumer electronics devices to the same network, particularly since the vast majority of homes do not have Category 5/6 Ethernet cabling and, in most cases, the MoCA network is already connected to the Internet. Although MoCA allows for password protection of the network, most service providers have not seen a need to enable this feature (content security is handled at higher layers). As a result, simple Ethernet to MoCA adapters can allow Ethernet connected consumer electronics devices like televisions, consumer STBs, gaming devices, etc., to easily connect to other devices in the home and to the Internet. Even if a home is not already MoCA

enabled by a service provider, multiple adapters can provide a simple connectivity solution to consumers looking to move video around their home.

An even more compelling use case is using MoCA as a backbone to support whole-home wireless coverage. Wi-Fi does not have ubiquitous high data-rate coverage in most homes, and this has been the primary barrier to service provider adoption of Wi-Fi for whole-home mDVR. With the rapid rise in the use of tablets and other wireless devices to receive video content, the need for whole-home wireless coverage has increased dramatically. Using MoCA to connect two or more Wi-Fi Access Points greatly enhances wireless coverage in the home. Together with more recent standards solutions such as IEEE 1905.1 (see Section V) that assist consumers in setting up complex networks, the consumer is now able to achieve a better quality of experience using their wireless video devices.

IV. TECHNOLOGY

A. The Home Coaxial Cable Infrastructure

A typical in-home coax architecture consists of a single coax point of entry (POE) at the side of the house fed through a drop cable from a multitap (directional coupler) in the street. The POE feeds outlets in multiple rooms through a series of splitters connected with RG-59 or RG-6 coax (see Fig. 2). The attenuation from the POE to each outlet is limited to about 25 dB in order to meet the receive level requirements of TVs and STBs.




Fig. 2. Home coax network example.

B. Channel Characteristics and Coexistence on the Coax

MoCA was designed to transport high-speed streaming video data over the existing coax cable distribution plants in homes. Coax cable is the ideal medium because television displays are usually already connected to the coax distribution system, and the coax cable is shielded and capable of carrying high bandwidth signals. The challenge of sending high-speed data over coax is that the distribution plants are not designed for room-to-room signal flow but are designed to feed a common signal to different rooms in the house. Coax distribution plants are typically designed with a trunk and branch topology where the coax outlet in each room is connected to the trunk through splitters and individual pieces of coax cable. The consequences of this topology are that room-to-room channel characteristics are not well controlled and may include severe multipath and high path loss. The MoCA system deals with these channels by using bitloaded orthogonal frequency-division multiplexing (OFDM) to adapt to each channel independently and by transmitting high signal power.

In order to operate on existing coax in the home with minimum disruption, MoCA was also designed to coexist with other signals on the coax. To avoid existing signals, MoCA is capable of operating on various frequencies. For example, in a home with CATV signals below 1 GHz, MoCA operates above 1125 MHz in a band called D. In a home with satellite L-band signals above 950 MHz, MoCA operates between 475 and 625 MHz in a band called E or between 650 and 875 MHz in the F-band (see Section IV-E6b). To further minimize interference to existing services, MoCA also features transmit power control (TPC).

C. MoCA 1.0 Field Trials

As mentioned above, a key requirement for service provider deployments is to maximize coverage in the home in order to minimize the remediation of problems which results in added cost to the service provider. This requirement can be quantified as throughput at greater than 95% coverage. Following the definition of MoCA 1.0, field trials were carried out in 250 homes in 15 cities across the United States. The primary goal was to determine realistic outlet coverage of the new standard. The key results measured and verified in the field tests were:

- PER less than $1\text{E-}6$;
- latency less than 5 ms;
- net usable data rates, with no changes to the home coax system:
 - 97% of all paths in all homes achieved >100 Mb/s (Fig. 3).

The results showed that the MoCA 1.0 technology supported key service provider requirements. The test results cited in this document are averaged across the 1000–1400 MHz band. The most important parameter for both the field testing and multimedia home networking




Fig. 3. Outlet coverage—percent of paths versus bitrate.

deployment is throughput, or net usable medium access control (MAC) data rate. The throughput and physical layer (PHY) level bit rates at the maximum transmission power as measured in the field tests are shown in the following curves.

D. Interference

There are alternate technologies that can operate on the low-frequency cable spectrum in the home (e.g., HPNA, HomePlug, and ITU G.hn on coax). Low radio frequency (RF) has the potential for greater reach, since cable attenuation is less than at high frequencies. This is countered by the fact that cable splitters leak much less at low frequencies and since the signal is always going “backwards” through one splitter (i.e., through the output ports), the higher leakage probability at high frequencies is actually of value. In addition, it is less well known that interference generated on the in-home power line couples onto the coaxial cable at low frequencies, resulting in significantly higher ingress at those frequencies. Fig. 4 shows the results of an experiment that measured the interference on a piece of coaxial cable that was placed next to a power cable. Various household appliances such as blenders and coffee grinders were plugged in to the power cable and turned on. A spectrum analyzer on max-hold captured the ingress on the coax. It is clear that powerline interference couples on to the coax until about 400 MHz. This is the reason that MoCA specified the lower edge of its frequency of operation at 400 MHz.




Fig. 4. Low-frequency ingress.

E. MoCA Technology

1) *Architecture*: The MoCA network is architected as a full-mesh or peer-to-peer network, as shown in Fig. 5. In this example, four STBs are connected to the in-home coax. These four devices result in six unique coax links (1 to 2, 1 to 3, 1 to 4, 2 to 3, 2 to 4, and 3 to 4). Because of differences in transmitter and receiver characteristics, a bidirectional MoCA link is established for each coax link resulting in 12 unique MoCA connections.

2) *Packet Type and Routing*: The MoCA network accepts and transparently bridges Ethernet packets. By providing support for Ethernet packets, MoCA can seamlessly integrate into any home data network. All data including video and audio must be encapsulated in Ethernet packets before being transported across the MoCA network. MoCA 1.x nodes use the MAC source address (SA), MAC destination address (DA) to route packets between MoCA nodes, and the virtual local area network (VLAN) priority (IEEE 802.1p) to prioritize the transmission of the packets across the MoCA network. MoCA 2.0 adds the option to use the differentiated services code point (DSCP) field in the IP header (if present) for prioritizing the packet transmissions. Because the MoCA network is not a broadcast network and packets are unicast between MoCA nodes, MoCA nodes cannot act as learning bridges by snooping traffic. Instead, each MoCA node monitors the SA of packets entering the MoCA network through it and distributes that information to all the other nodes via notification messages. The notification messages inform other nodes that any messages destined to that address should be forwarded to the notifying node. The whole MoCA network behaves like a distributed Ethernet switch with QoS, where each MoCA node works like a port of the Ethernet switch, and notification messages between MoCA nodes are like communications between Ethernet ports of the switch.

3) *Throughput, Latency, and Packet Loss*: The MoCA system was designed to deliver “no excuses” streaming

video around the house. To meet this goal, MoCA targeted high throughput, low latency, and low packet loss. High throughput was required to accommodate transporting multiple simultaneous high-definition television (HDTV) streams. A single HDTV stream could typically range anywhere from 4 to 20 Mb/s, while trick modes (e.g., fast forward or reverse) might peak at two to four times higher rates. A home with four HD streams could consume 80 Mb/s plus require even more for trick modes. A MoCA 1.1 network can deliver more than 175 Mb/s of network throughput, and MoCA 2.0 increases this to over 400 Mb/s.

Low latency is important for various reasons: to minimize video delay and buffering; to allow digital rights management (DRM) schemes such as digital transmission content protection over IP (DTCP-IP) to work; to maximize transmission control protocol over IP (TCP-IP) performance; and to enable high-performance gaming. MoCA 1.1 has a maximum latency of 7.5 ms. Besides introducing delay, latency adds cost for buffer memory, since packets “in-flight” across the network must be buffered in data buffers somewhere in the MoCA network. To accommodate a 7.5-ms delay, the MoCA devices would have to buffer up to 7.5-ms worth of data: at 100 Mb/s, a 750-kb buffer is required. As delay and throughput increases, the size and cost of the buffers go up. DTCP-IP is one of the DRM schemes adopted by the Digital Living Network Alliance (DLNA). To operate, DTCP-IP monitors round-trip latency between the server and the client to ensure they are both on the same local network and not separated by the Internet cloud. If latency is too large DTCP-IP will shut down the flow. Hence, low latency is required to support DLNA networks. Low latency is also needed for best TCP-IP performance. It is well known that TCP-IP throughput is inversely proportional to the round-trip delay [3], so as delay increases, TCP-IP slows down. Tests have shown that latency in MoCA 1.x can limit TCP-IP performance to less than 100 Mb/s. One of the design goals for MoCA 2.0 was to improve TCP-IP performance by reducing latency. As shown in Table 2, the average latency for a MoCA 2.0 network was reduced to 3.6 ms compared to




Fig. 5. MoCA network architecture.

4.5 ms for a MoCA 1.1 network. In addition, special support for gaming was added by enabling a low-latency mode guaranteed to have average latency less than 2.4 ms and maximum latency less than 3 ms.

For streaming video (i.e., IP multicast), where there is no data retransmission, low packet loss is required to deliver high-quality glitch-free video. MoCA 1.x guarantees a packet loss rate of less than $1e-5$, and MoCA 2.0 lowered the allowed packet loss rate to less than $1e-6$ or $1e-8$ selectable by the higher layers.

4) *Security*: MoCA provides a privacy feature. When it is turned on, all data packets are encrypted. MoCA 1.x uses 56-b Data Encryption Standard (DES) encryption, while MoCA 2.0 upgraded the encryption to 128-b Advanced Encryption Standard (AES). Every MoCA device has a programmable password. Devices which share the same password will form a network with each other, while devices with different passwords will avoid each other. Passwords are useful for distinguishing between networks in adjacent homes where the MoCA signal can leak between homes over a cable operator's network.

5) *Quality of Service*: In addition to streaming video, the MoCA network may be used to simultaneously transport other services such as best effort Internet traffic, file transfers, etc. When this happens, the video stream will contend for bandwidth on the MoCA network, which may result in dropped video packets. MoCA offers two methods to protect the video traffic from other traffic. The first method is to send video as prioritized traffic. MoCA inspects packets which have VLAN tags and prioritizes traffic based on the three priority bits in the VLAN tag. By sending video with high VLAN priorities, MoCA will give video precedence over untagged and lower priority packets and eliminate any contention between them. The second method to protect the video is to send video using parameterized quality of service (PQoS). Within MoCA, traffic flows with specific flow parameters can be admitted to the network. A traffic specification (TSPEC) is used to configure the key metrics for a flow (see Table 1). Once admitted, these flows are guaranteed the bandwidth required to be transported across the network. PQoS

flows can be uniquely identified by the content of their packet headers. Packets classified as belonging to a PQoS flow are assigned the highest priority, and their delivery is guaranteed. A unique aspect of the MoCA PQoS architecture is the ability of lower priority traffic to make use of reserved bandwidth that is not being used at that time (this is typical of variable bit rate video traffic which is bursty in nature). MoCA 2.0 adds additional TSPEC parameters for greater control (e.g., maximum latency, classification rule, retransmission, in-order packet delivery, PER).

6) *MoCA 1.x Technical Background*: Table 2 lists basic parameters of MoCA 1.0, MoCA 1.1, and MoCA 2.0. This section provides more information about MoCA 1.x, while in Section IV-E8 will explain changes made for MoCA 2.0.

a) *MoCA 1.x PHY Layer*: The MoCA 1.0 and MoCA 1.1 PHY layers are identical. Both use bitloaded OFDM with 224 subcarriers in a 50-MHz channel. Bit-loaded OFDM was selected for MoCA because it is robust against static or slowly changing multipath and optimizes the modulation between every pair of devices. An example of bitloading is shown in Figs. 6 and 7. The individual OFDM subcarriers are loaded with quadrature amplitude modulation (QAM) constellations that support a required PER.

The MoCA 1.0 field trial showed that this technique has a 95% chance of delivering 100 Mb/s between any two rooms in a home without any remediation. This translates to 140 Mb/s for MoCA 1.1. Bit loading is performed periodically between all the nodes to track slow changes in hardware or the cable distribution plant due to temperature or other effects. When bit loading, each MoCA device probes the channel between itself and every other MoCA device in the network and selects the modulation on each of the 224 subcarriers based on the probe results: the better the signal-to-noise ratio (SNR) on a subcarrier, the higher the modulation assigned to that subcarrier. MoCA 1.x uses a maximum subcarrier modulation of 256 QAM. Since the MoCA PHY layer adapts each link between node pairs independently, the channel capacity can be different between different nodes as well as between the forward and reverse directions of the same nodes. The bit-loading parameters for a particular path are called a PHY profile. Fig. 8 illustrates how the MoCA network is composed of a

Table 1 MoCA 1.1 PQoS Traffic Spec Parameters

MoCA TSPEC Parameter	Description
Peak Data Rate	The maximum amount of payload per second that a PQoS flow transfers over the MoCA network.
Packet Size	The packet length of a flow which is used to calculate the bandwidth cost of the flow.
Lease Time	The amount of time until the nodes automatically release resources for the flow. A flow can be configured for an infinite lease.

Table 2 MoCA Comparison Chart

	MoCA 1.0	MoCA 1.1	MoCA 2.0	MoCA 2.0 Bonded
Data Modulation	OFDM	OFDM	OFDM	OFDM
Reservation Request Modulation	OFDM	OFDM	OFDM and OFDMA	OFDM and OFDMA
Number of subcarriers	224	224	480	960
Subcarrier Spacing	195.313 kHz	195.313 kHz	195.313 kHz	195.313 kHz
Max Subcarrier Modulation	256 QAM	256 QAM	1024 QAM	1024 QAM
Channel BW (MHz)	50	50	100	2x100
Forward Error Correction	Reed Solomon	Reed Solomon	LDPC	LDPC
Encryption	56 bit DES	56 bit DES	128 bit AES	128 bit AES
Maximum Number of network Nodes	8	16	16	16
Max two node Throughput	>125 Mbps	>176 Mbps	>400 Mbps	>800 Mbps
Max Throughput > 6 nodes	>98 Mbps	>140 Mbps	>400 Mbps	> 800 Mbps
Turbo Mode Throughput (2 nodes only)	NA	NA	>500 Mbps	>1 Gbps
Prioritized QoS	Yes	Yes	Yes	Yes
Parameterized QoS	No	Yes	Yes	Yes
Maximum Packet Loss	<1e-5	<1e-5	<1e-6 (normal PER mode) <1e-8 (Very low PER mode)	<1e-6 (normal PER mode) <1e-8 (Very low PER mode)
Backwards Compatibility	Baseline	MoCA 1.0	MoCA 1.0, MoCA 1.1	MoCA 1.0, MoCA 1.1
Latency	<3.5 ms (avg) < 5 ms (95%)	< 4.5 ms (avg) < 7.5 ms (99.99%)	<3.6 (avg) <5 .1ms (99.9%)	<3.6 (avg) <5.1 ms (99.9%)
Transmit Power Control	Yes	Yes	Yes	Yes
Retransmission	No	No	Yes	Yes
Fragmentation	No	No	Yes	Yes
Low Power Modes	No	No	Yes	Yes
Ranging	No	No	Yes	Yes
Multicast Support	Broadcast	Broadcast	Multicast Flow	Multicast Flow

collection of individual PHY profiles which can have different channel capacities. Whenever data are sent between two MoCA nodes, the PHY profile between those two nodes is used.

In addition to unicast PHY profiles between node pairs, MoCA nodes also create a broadcast PHY profile between each node and all the other nodes in the network. The broadcast PHY profile is used to transmit broadcast and multicast data. This profile is constructed from the greatest

common density (GCD) of the individual subcarrier constellations (see Fig. 9) and allows all nodes to “hear” the transmitting node. This mode of operation is less efficient than a single unicast transmission but considerably more efficient than multiple unicast transmissions of the same information.

MoCA PHY layer packets are sent as shown in Fig. 10. Packets start with an interframe gap (IFG) where nothing is sent. The IFG separates the packet from the previous




Fig. 6. Bit-loaded OFDM example—“good channel.”




Fig. 7. Bit-loaded OFDM example—“impaired channel.”

packet. The IFG is followed by a preamble which is used to synchronize to the start of the packet. Channel estimation (CE) symbols are then sent, followed by the MAC frame (see Section IV-E6c) and associated forward error correction [MoCA 1.x packets include Reed–Solomon (RS) parity bits in the data symbols to help keep packet loss below $1e-5$ —a level acceptable for video]. Finally, the packet ends with a cyclic redundancy checksum (CRC).

Since MoCA shares the coax with other signals, TPC is used between MoCA nodes to minimize interference to other devices using the coax. During each bit-loading cycle, the signal power between nodes is adjusted along with the modulation. TPC can reduce transmit power by up to 30 dB.

b) *MoCA 1.x Channel Plan*: The MoCA PHY layer can also operate at different frequencies. Fig. 11 shows the frequency plan for MoCA 1.1. Table 3 shows the range of

center frequencies that can be used by MoCA 1.x devices. The actual MoCA 1.x channel occupies 50 MHz. D-band is used when sharing the coax with a cable or telco operator signal and E- and F-bands are used when sharing the coax with a satellite signal. Bands A and B are currently not used and C-band is used for the WAN link in Verizon FiOS home, as described in Section III-A.

The MoCA specification defines channel scanning requirements to allow a new node to find an existing network (for example, in the D-band) and ensure that multiple networks are not formed. In order to maintain consistent and robust behavior after power failures, a node records the frequency of the network as the last operational frequency (LOF) and uses that frequency during future network searches.

c) *The MoCA 1.x MAC Layer*:

MAC Packet Encapsulation: The process used to create a MAC frame payload from an Ethernet packet is illustrated in Fig. 12. As shown, the frame checksum of an Ethernet packet is dropped at the source end. The




Fig. 8. Channel capacity between nodes.




Fig. 9. Broadcast bit loading.

IFG	Preamble	CE Symbols	MAC Frame	FEC	CRC
-----	----------	------------	-----------	-----	-----

Fig. 10. PHY layer packet.

destination address contained in the Ethernet packet is used to derive destination node IDs of the MAC frame header for the encapsulated packet. A payload CRC is then added.

MAC Scheduling: The MoCA MAC uses a time-division duplex (TDD) scheme where all nodes transmit on the same frequency during different time slots. The time slots are coordinated by a single node called the network coordinator (NC). The NC is dynamically selected from all the nodes in a network based on best broadcast bit loading. To ensure all nodes operate on the same time clock, all the nodes synchronize their clocks to the NC's time, which is broadcast to all the nodes approximately every millisecond in what is known as a MAP message. MAP messages also carry dynamic scheduling information which defines when each node can transmit in an upcoming time period called a MAP cycle (literally, a map of the scheduling cycle). By specifying how every time slot is allocated, the NC ensures no contention or collisions occur on the network which is necessary to support high reliability needed for streaming video.

During each MAP cycle, some MoCA nodes are given the opportunity to send a reservation request (RR) message to the NC. In their RRs, nodes request time slots to send messages to other nodes. The NC responds to all the RRs it receives in a MAP cycle by granting time in the next MAP cycle for as many transmissions it can and sending a schedule for those grants in the next MAP message. When nodes receive the schedule in the MAP message, they know when they should send and receive data during the upcoming MAP cycle. Nodes which have received Ethernet packets from an upper layer to send across the MoCA network are called ingress nodes. In the RR requests, ingress nodes include priority information and MoCA 1.1 nodes also include PQoS information. The NC uses this information to schedule PQoS data first followed by prioritized data in priority order followed by

best effort data. Since the NC sees the priorities of all the packets in the network, it is able to prioritize packets on a network-wide basis and not just locally. If nodes request more data than the network can carry, lowest priority packets are dropped first from the schedule.

Table 2 shows that MoCA 1.x data throughput declines as a MoCA network goes from two nodes to more nodes. The decline is a direct result of the overhead associated with granting additional nodes opportunities to send RRs. As new nodes are added to a network, the NC must schedule additional RRs which take away transmission time slots from data. The NC will add additional RRs per MAP cycle for each node up to six nodes, including the NC, after this the NC will use "multicycle" scheduling where nodes do not get an RR grant every MAP cycle. Instead they may only get an RR opportunity every two or three MAP cycles.

Latency in the MoCA network is also dominated by this request-grant process. From the time an Ethernet packet arrives at an ingress node, to the time the node gets a grant time to send the packet several milliseconds may have passed. Total latency comprises the following times.

- 1) Packet arrival to RR time: The time between the arrival of an Ethernet packet at the ingress node until the time the node can transmit an RR which requests a transmission slot for that packet. This may take several MAP cycles.
- 2) RR to MAP time: The time between an ingress node sending an RR and receiving the next MAP which grants time slots to transmit.
- 3) MAP to grant time: The time from receiving the MAP until the actual transmission time granted in the MAP.
- 4) Egress time: The time it takes for the receiving node (egress node) to forward the packet after receiving the packet on the coax.

Although MoCA 1.0 and MoCA 1.1 use the same request-grant mechanism and the same MAP cycle




Fig. 11. MoCA 1.1 channel plan.

Table 3 MoCA 1.1 Channel Frequencies

Band	Center Frequency (MHz)
E	500-600
F	675 - 850
A	975
B	900
C	925 - 1000
D	1150 - 1500

duration, a MoCA 1.1 network can have more latency than a MoCA 1.0 network (see Table 2). The reason for this difference is because a MoCA 1.1 network can have 16 nodes while a MoCA 1.0 network is limited to eight nodes. In MoCA 1.x, the NC typically grants five or less RR opportunities during each MAP cycle so when there are more than six nodes in the network (including the NC), a node must wait longer between RR opportunities. So since there may be more nodes in a MoCA 1.1 network, the MoCA 1.1 nodes may experience longer packet arrival to RR times. Latency in a MoCA 1.1 network with eight nodes or less will be the same as latency in a MoCA 1.0 network with the same number of nodes.

To achieve higher throughput, MoCA 1.1 adds packet aggregation to MoCA 1.0. Packets from an ingress node to the same egress node can be aggregated together and sent in one PHY layer packet (Fig. 10). By aggregating packets, the PHY layer overhead is shared and MoCA 1.1 can get a maximum throughput of 176 Mb/s, compared to 125 Mb/s for MoCA 1.0. MoCA 1.1 allows aggregation of up to 6 kB into one PHY layer packet.

d) *Management and Diagnostics*: Since service providers deploy almost all MoCA networks as part of a paid service, it is of critical importance to ensure a good user experience from the date of installation onwards. To assist in this, MoCA provides an extensive set of management and diagnostic capabilities that can be accessed locally during installation or remotely. This information is specified within a management information base (MIB) framework to leverage common SNMP protocols as well as through the Broadband Forum's TR-069 specification. Table 4 gives some examples of the available parameters.

7) *Service Provider Deployment Considerations*: Cable service providers typically install a POE filter prior to the POE on the side of the house. This is a low-pass filter that prevents MoCA signals from leaking onto the external coax plant. It has an added benefit of reflecting signals inside the house, thus improving the outlet coverage even further.

A small number of homes have in-line amplifiers in the house to boost the cable TV signal in the cable band. These amplifiers can severely attenuate the MoCA signal above 1 GHz and require remediation by the cable operator during the installation process. Remediation can include replacing the amplifier with a MoCA-ready amp that allows the MoCA band to bypass the amplifier, or installing a new amplifier prior to the first splitter, thus allowing communications on the coax in the home to be unaffected.

8) MoCA 2.0

a) *Throughput Improvements*: A major goal of MoCA 2.0 over MoCA 1.1 was to increase overall throughput and improve TCP-IP performance while maintaining backwards compatibility. MoCA 2.0 increases 16 node throughput from > 140 Mb/s in MoCA 1.1 to > 400 Mb/s.

At the PHY layer, the following changes were made in MoCA 2.0 to increase throughput:

- increased channel bandwidth (BW) from 50 to 100 MHz;
- increased number of OFDM subcarriers from 224 to 480;
- increased maximum modulation from 256 to 1024 QAM per subcarrier;
- forward error correction (FEC) changed from RS to low density parity check (LDPC);
- reduced IFG;
- ranging;
- OFDMA for reservation requests.

Increasing the channel bandwidth from 50 to 100 MHz and increasing the number of subcarriers from 224 to 480 more than doubled the number of bits that could be sent per OFDM symbol. By increasing the maximum modulation from 256 to 1024 QAM and changing the FEC from RS to LDPC, another 10% gain was achieved. Although LDPC decoding is significantly more complex than RS decoding, it has more coding gain which allows 1024 QAM


**Fig. 12.** MAC frame encapsulation.

Table 4 Management Parameters

Parameter	Description
RF Channel	Reports RF channel for which the node is tuned
Last Operational Frequency	Last known operating frequency of the node
Taboo List Start Freq	Defines the position of the active channel in the Taboo Mask
Taboo List Bit Mask	Defines off-limits channels above and below the active channel
Modulation Profiles	Reports the unicast modulation profiles for each connected node
Transmit (Tx) Broadcast profile from this node to all other nodes	Reports the Broadcast profile for the Tx of the node being queried
Estimated Receive (Rx) level in dBm	Estimation of Rx level per modulation profile
Tx Power Control Reduction	Tx power reduction from this node to other nodes
Power Control Target PHY Rate	Target PHY rate used for transmit power control
MAC Address	Device MoCA interface MAC address
MoCA Version	MoCA version supported
MoCA Network Status	Reports whether node is “up” (e.g. part of a MoCA network) or searching for a network
Connected Nodes	Reports the number of nodes on the network
Privacy Status	Reports whether Modulation Privacy is enabled or disabled
Privacy Password	Privacy Password
Packets Transmitted on the MoCA interface	Cumulative count of Tx data packets on the MoCA interface since reset of the interface
Packets Received on the MoCA interface	Cumulative count of Rx data packets on the MoCA interface since reset of the interface
Missed packets or packets with CRC errors per node	Total number of packets that were not detected by the node or packets that had CRC errors. This count is accumulated per each transmitting node
Count of corrected Reed-Solomon bytes	Cumulative count of corrected R-S bytes per Rx modulation profile channel

to be received with only 1 dB more SNR compared to 256 QAM with RS coding.

Ranging adjusts the transmit time of packets so that the packets arrive at a receiver when they are expected. By adding ranging to MoCA 2.0, transit time and timing errors between nodes are compensated thus reducing the variability of packet arrival times. This enables using smaller IFGs. The smaller IFG has a minor impact on maximum throughput with large packets but has a bigger impact on small packets. By looking at Fig. 10, one can see how, for large packets with many data symbols, the IFG will be a small portion of the total packet size but, for small packets with a few data symbols, the IFG will be a large portion of the total packet. Along the same lines, as we increased the number of bits per OFDM symbol, the PHY packets get shorter because it takes fewer data symbols to send the same data, thereby decreasing PHY efficiency. To mitigate this, MoCA 2.0 also increases the aggregation size from 6 to 12 kB, which offsets the loss due to the higher bits per OFDM symbol.

A major improvement in MoCA 2.0 was to use orthogonal frequency division multiple access (OFDMA) for RRs. RRs are, by nature, short packets, and in MoCA 1.x, RRs incur a large percentage of overhead in their IFG, preamble, and CE symbols shown in Fig. 10. With OFDMA, eight nodes simultaneously send their RRs to the NC by each transmitting on a different set of subcarriers. By

using OFDMA, eight nodes can share the same IFG and CE symbols, and no preamble is necessary. In comparison, a 16-node MoCA 1.x network uses approximately 25% of a MAP cycle to transmit five RRs, and it takes three MAP cycles to grant all 15 non-NC nodes an RR while with OFDMA in a MoCA 2.0 network 15 RRs every MAP cycle take approximately 8% of the MAP cycle. Not only does this reduce the overhead for RRs, but also it reduces latency by allowing MoCA 2.0 NCs to grant RR opportunities to all nodes every MAP cycle. This reduction in latency translates into less data buffering in the network and reduces memory cost which offsets the need for more memory due to the higher throughput.

In MoCA 1.1, when the transmissions are distributed over multiple nodes, the total network throughput is reduced to the MoCA 1.0 level due to not using aggregation. MoCA 2.0 added latency tolerant PQoS flows which can specify latencies beyond 10 ms that can be accommodated. Flows of this nature will be able to hold and aggregate packets at the ingress node and by doing so enable full MAC throughput, even when the transmissions are distributed among multiple nodes

b) *MoCA 2.0 Channel Plan*: MoCA 2.0 occupies a 100-MHz channel, so there are fewer nonoverlapping MoCA channels available to the MoCA 2.0 devices in the D-band. To add flexibility to frequency planning, MoCA




Fig. 13. Single or bonded channel in the extended D-band.

extends the D-band from 1525 to 1675 MHz and introduces a split D-band (subband D-low and subband D-high) so two independent MoCA networks can be supported. MoCA 2.0 single channels (100-MHz BW) are centered on a 25-MHz grid and can tune in 25-MHz increments. Bonded-pair channels (225-MHz BW) have a fixed 25-MHz gap between them where both the primary channel and secondary channels are centered on the 25-MHz grid. This frequency plan supports any of the following combinations in the extended D-band:

- single or bonded-pair MoCA 2.0 channels;
- one MoCA 2.0 channel plus one independent single or bonded-pair MoCA 2.0 channel;
- mixed mode operation, comprising a mix of MoCA 1.x and 2.0 nodes.

These possibilities are illustrated in Figs. 13 and 14.

The E-band in MoCA 2.0 is defined between frequencies 400 and 700 MHz (300 MHz wide). This frequency plan enables:

- MoCA 2.0 single-channel or bonded-pair-channel operation in the E-band;
- mixed-mode operation (MoCA 1.x/MoCA 2.0) in the E-band.

Fig. 15 illustrates the MoCA 2.0 frequency plan in the E-band with the bonded-pair-channel example.

The F-band is defined between frequencies 650 and 875 MHz (225 MHz wide). This frequency plan enables single or bonded-pair operation and mixed-mode operation and is illustrated in Fig. 16.

MoCA defines a strict set of coexistence requirements for each of the above bands as well out-of-band interference specifications to support coexistence with separate services on the same coaxial cable.

c) *Power Consumption and Regulation:* Energy efficiency has become a mainstream requirement for all consumer products and specifically STBs and networking equipment. U.S. federal and state governments as well as the European Commission have initiatives to regulate the maximum allowed annual energy consumption of STBs and networking equipment [4]–[7].

Given the expected power consumption of STBs and other network devices with MoCA, a practical way to limit the annual energy consumption is to minimize power consumption during times when these devices do not need to be fully functional. In support of the effort to reduce annual energy consumption, MoCA 2.0 adds support for




Fig. 14. Two independent networks in the extended D-band.




Fig. 15. MoCA 2.0 E-band frequency plan with the single-channel or bonded-pair-channel example.




Fig. 16. MoCA 2.0 F-band frequency plan with the bonded-pair-channel example.

multiple power states, allowing a MoCA device, under the control of its host processor, to move into and out of low power states in coordination with the other devices in the network. By doing so, MoCA 2.0 can reduce the annual energy consumption of the MoCA home network as well as enable significantly lower energy consumption of all devices connected on that network such as Gateways, Media Servers, IP-STB, and MoCA-Wi-Fi extenders. The MoCA 2.0 power modes are shown in Table 5. For backwards compatibility, when a MoCA 1.x device is in the network, all the MoCA 2.0 devices enter the M0 power mode. A description of the MoCA 2.0 low-power modes is provided in Table 5.

MoCA 2.0 defines an interface between the node and a higher layer power management subsystem (PMS) and specifies the rules for transitioning the node from the active state to any of the other power states and from the other power states back to the active state.

In a home network comprising a media server and multiple thin clients, all connected through the MoCA network, the different nodes may use different power modes

to save energy. Fig. 17 shows two examples of devices transitioning to different power states. A thin client which does not record any broadcast service can go into a lower power mode based solely on user input. As an example when the user turns off the TV and pushes the standby (STBY) button of the thin client, the node in the thin client can move to the sleep state (M3). To move to the sleep state, the node requests and gets a grant from the NC for the move and then can lower its power consumption and be in the sleep state. This node needs only to wake up when the user starts watching TV again. When that happens, the PMS in the thin client will instruct the MoCA node to wake up and the node will rejoin the MoCA network and resume full operation. Fig. 17 shows an example of an IP-STB going into sleep state following a user pushing the STBY button of the remote control, and out of sleep state following a user pushing the ON button of the remote control. The media server, on the other hand, may receive a request at any time from any of the client nodes in the MoCA network to start forwarding a recorded movie. To be able to send such a request to the media server through the MoCA network while the media server is in a low-power state, the node in the thin client has to wake the media server up. This can only be done if the node in the media server limits its low power states to either idle (M1) or standby (M2). These are the power states that support wake-on-MoCA (WoM). As shown in Fig. 17, when the media server is in standby (M2), as soon as the thin client (the MoCA Wi-Fi extender) gets a request to start watching a movie it sends a request to the NC to wake the media server up and the NC sends a WoM message to the node in the media server. At that point, the node in the media server rejoins the network and the media server returns to full

Table 5 MoCA 2.0 Low-Power Modes

Power Mode	Name	Description
M0	Active	Normal operation. Full power consumption.
M1	Idle	Node can receive broadcast and unicast traffic. Node does not transmit data traffic. Node wakes-up on MoCA traffic. Fast wake-up.
M2	Standby	Node can receive broadcast packets. Node does not transmit data traffic. Node wakes-up on MoCA traffic. Slower wake-up.
M3	Sleep	Node is disconnected from the network.




Fig. 17. Power state transition examples in a MoCA home network.

operation in order to send the recorded movie to the client node. The decision of the media server on whether to use idle or standby is a tradeoff between the requirements of power savings and quickness of waking up the node from the low-power state.

MoCA Nodes Energy Consumption: Table 6 shows an example of the potential savings in the yearly energy consumption between a MoCA 1.x node and a MoCA 2.0 node, assuming, for example, that MoCA 1.x always consumes 2 W and MoCA 2.0 consumes 4 W in the active state, 1 W in the idle state, and 0.3 W in the standby/sleep state. The MoCA 1.x node is assumed to always be in the active position, and the MoCA 2.0 node is assumed to be in one of the low-power states for 17 h each day.

Table 6 MoCA Yearly Energy Consumption

Use Case	Yearly Energy Consumption
MoCA 1.x 100% active	$2 \times 24 \times 365 = 17.5$ KWh/year
MoCA 2.0 7h Active, 17 h Idle (per day)	$(4 \times 7 + 1 \times 17) \times 365 = 16.4$ KWh/year
MoCA 2.0 7h Active, 17 h Standby/Sleep (per day)	$(4 \times 7 + 0.3 \times 17) \times 365 = 12.1$ KWh/year

Table 7 STB Yearly Energy Consumption

Use Case	Power Consumption in Low Power	Yearly Energy Consumption
Current STB	N/A	$20 \times 24 \times 365 = 175$ KWh/year
MoCA 2.0 based STB	5W	$(20 \times 7 + 5 \times 17) \times 365 = 82$ KWh/year

Table 7 shows an example of the potential energy savings of the new generation STBs. In this example, the power consumption of the STB is assumed to be 20 W when it is fully functional (active state) and 5 W when it is in low-power state.

d) **Other Changes:** In addition to the above changes, MoCA 2.0 added the following enhancements:

- turbo mode for two nodes only;
- bonded mode for higher throughput;
- fragmentation;
- retransmission;
- 128-b AES encryption;
- multicast;
- policing of PQoS flows;

- support for 802.1 audio/video bridging (AVB);
- support for multiple networks.

In certain MoCA applications, there may only be the need to support two nodes. To optimize these cases, MoCA 2.0 includes a turbo mode which supports only two nodes. In this mode, the nodes may eliminate some MAC overhead and be more efficient so they are required to achieve a higher maximum throughput between them. The maximum MAC throughput in turbo mode is required to be > 500 Mb/s.

MoCA 2.0 also introduces a channel bonded mode. In this mode, two physical MoCA channels are used to create one logical channel. The two physical channels have a 25-MHz guard band between them and occupy a total of 225 MHz. When transmitting between two bonded MoCA 2.0 devices, the nodes will transmit half the data on each channel simultaneously. This doubles the achievable MAC throughput from > 400 to > 800 Mb/s and from > 500 Mb/s to > 1 Gb/s for turbo mode.

Fragmentation gives the NC a degree of freedom which allows it to be more efficient in scheduling packets in a MAP cycle. MoCA 1.x has scheduling rules which often result in the MoCA 1.x NC wasting time on the wire when it cannot schedule packets. Fragmentation allows the NC to break packets into smaller ones, thus filling the time gaps which MoCA 1.x had to waste. The benefit of fragmentation varies depending on traffic flow but fragmentation can increase throughput by $> 7\%$. Fragmentation is only enabled between MoCA 2.0 nodes.

Although the coax environment is generally static and bitloading matches the environment well, there are cases when transient events may cause packet loss. Transients may be introduced by momentary ingress into the coax or quick changes in coax topology such as connecting or disconnecting a device from the coax. Retransmission has been added to MoCA 2.0 as a tool to combat situations where small data loss cannot be tolerated. MoCA 2.0 allows retransmission to be turned on for a unicast flow between two nodes. Retransmission is selective and packets which are retransmitted will experience some extra delay.

AES is the newest NIST encryption standard. For higher security, MoCA 2.0 upgraded its cryptoalgorithm to 128-b AES from 56-b DES in MoCA 1.x. MoCA 1.x nodes include a user entered password which is used to generate static 56-b DES keys. MoCA 2.0 devices use the same password to generate static 128-b AES keys. For backwards compatibility, the same password is used in MoCA 2.0 and MoCA 1.x modes.

Support for multicast flows over the MoCA network allows the receiving devices to select (e.g., based on their multicast memberships) which flows to forward to their Ethernet convergence layer and to drop without processing, and by doing that save processing power both at the MoCA level and at the host which would otherwise receive these flows. Multicast flows are supported in both prioritized flows and PQoS flows.

The policing of PQoS flows provides a tool which can be used to either enforce the PQoS requirements on the ingress nodes or as a debugging tool to identify flows which violate their admitted PQoS flow parameters.

The IEEE 802.1 AVB Task Group provides specifications that allow time-synchronized low latency streaming services through 802 LANs. MoCA 2.0 specifies the requirements that a MoCA device has to meet in order to support AVB together with other MoCA devices in the network. These requirements include both the time stamping of AVB messages and the encapsulating of AVB messages in special MoCA messages.

MoCA 2.0 requirements are designed to allow the coexistence of more than one MoCA 2.0 network in the home. A MoCA network in the D-band (or D-low or D-high) can coexist with another MoCA network in either the E-band or the F-band. A MoCA network in band D-low can also coexist with another MoCA network in D-high, and at the same time with a network in the E-band or the F-band.

V. RELATED STANDARDS

Some of the other standards that make up part of the in-home multimedia network story are described below.

DLNA defines interoperability guidelines to enable sharing of digital media between multimedia devices. DLNA is built on top of existing specifications but specifies various restrictions on those standards to allow for true interoperability. DLNA allows devices to discover each other and their digital content and stream that content between devices. Device classes specified by DLNA include such elements as: digital media server (DMS), digital media renderer (DMR), digital media player (DMP) and digital media controller (DMC). MoCA is one of the approved networking technologies in DLNA. DLNA also supports a commercial video profile (CVP) which allows service providers to securely distribute their protected media to consumer electronics devices in the home.

IEEE 1905.1 is a standard that specifies an abstraction layer for multiple home networking technologies (including IEEE 802.3, IEEE P1901, IEEE 802.11, and MoCA 1.1). The abstraction layer provides a common data and control service access point to these technologies. The standard supports dynamic interface selection and simplified user experience for adding devices to the network, setting up encryption keys, and extending network coverage. It also provides network management features related to network topology, path selection, and network control.

The Broadband Forum TR-069 specification defines an application layer protocol for remote management of end-user devices such as modems, routers, gateways, STBs, and VoIP phones. Through TR-069, service providers can remotely monitor, diagnose, and manage devices on the in-home network. The network itself is able to be managed

through data models that are unique to the specific home network. MoCA is one of the networking technologies supported by the standard.

VI. CONCLUSION

This paper has presented an overview of the MoCA market and technology. In Section II, we discussed the history behind the early drivers of a multimedia home network

and the genesis of MoCA. Section III briefly explored the service provider and consumer markets related to MoCA. Section IV presented the technology of MoCA, including the home coax architecture and the MoCA specifications. Finally, Section V presented information on related standards. The MoCA specification provides a solution to the high-speed and robust networking requirements on existing home cabling that has been driven by service providers for in-home video distribution. ■

REFERENCES

- [1] *MoCA 1.1 Specification*, Multimedia over Coax Alliance, 2005. [Online]. Available: www.mocalliance.org
- [2] *MoCA 2.0 Specification*, Multimedia over Coax Alliance, 2010. [Online]. Available: www.mocalliance.org
- [3] V. Jacobson, R. Braden, and D. Borman, *TCP extensions for high performance*, RFC 1323, May 1992.
- [4] European Commission, ICT Codes of Conduct. [Online]. Available: http://re.jrc.ec.europa.eu/energyefficiency/html/standby_initiative_main.htm
- [5] Regulations.gov, Energy conservation program for consumer products and certain commercial and industrial equipment: Proposed determination of set-top boxes and network equipment as a covered consumer product. [Online]. Available: <http://www.regulations.gov/#!docketDetail;dc=FR+PR+N+O+SR+PS;rpp=10;po=0;D=EERE-2010-BT-DET-0040>
- [6] Energy Star. [Online]. Available: http://www.energystar.gov/index.cfm?fuseaction=find_a_product.showProductGroup&pgw_code=ST
- [7] California Energy Commission. [Online]. Available: <http://www.energy.ca.gov/appliances/2011rulemaking/index.html>

ABOUT THE AUTHORS

Anton Monk received the M.S. degree in electrical engineering from the California Institute of Technology, Pasadena, CA, USA, in 1991 and the Ph.D. degree in electrical engineering from the University of California San Diego, La Jolla, CA, USA, in 1994.

He is the cofounder and V.P. Technology at Entropic Communications, Inc., San Diego, CA, USA, where he is responsible for new technology initiatives and standard activities. Previously, he led the engineering development of Entropic's coax-based multimedia networking chipset and technology which is now the *de facto* multimedia networking standard for service providers in the United States. He is CTO of the Multimedia over Coax Alliance and serves on the Multimedia over Coax Alliance (MoCA) Board of Directors and the RVU Alliance Board of Directors. Prior to cofounding Entropic Communications, Inc., in 2001, he was involved in the development of cable and satellite integrated circuits at Conexant, Irvine, CA, USA and ComStream Corporation, San Diego, CA, USA, and was a Researcher at the Jet Propulsion Laboratory, Pasadena, CA, USA.



Yoav Hebron received the B.Sc. and M.Sc. degrees in electrical engineering from the Technion—Israel Institute of Technology, Haifa, Israel, in 1978 and 1985, respectively.

He is Senior Director of Technology at Entropic Communications, Inc., San Diego, CA, USA. He is currently responsible for all technical aspects of the Multimedia over Coax Alliance (MoCA) standard activities within Entropic Communications, Inc. He is the Chair of the MoCA Specifications Work Groups and the editor of the MoCA 1.1 specifications and the Certification Test Procedures. From 1997 to 2007, he was Division Director of Systems Engineering at Conexant, Irvine, CA, USA, where he managed cable modem projects. From 1991 to 1997, he was Director of Systems Engineering at ComStream Corporation, San Diego, CA, USA, where he was involved in the development of cable and satellite communication systems and integrated circuits. From 1983 to 1990, he was an R&D Engineer at the Department of Communications, Rafael Armament Development Authority in Israel, where he was involved in research and development in digital communications. From 1978 to 1983, he served in the Israeli Defense Forces (IDF) as a Communication Engineer Officer.

Ronald Lee received the B.S. degree in electrical engineering from the Massachusetts Institute of Technology, Cambridge, MA, USA, in 1975 and the M.S.E.E. degree from the University of Illinois at Urbana-Champaign, Urbana, IL, USA, in 1976.

He is the Director of Strategic Technology at Entropic Communications, Inc., San Diego, CA, USA. Before joining Entropic Communications, Inc., he was Director of Systems and Technology in the Cable Access group at 3Com, Chicago, IL, USA, where he was responsible for cable modem development and technology. He was also Manager of R&D at Zenith Electronics, Chicago, IL, USA, where he worked on the ATSC HDTV transmission standard.