The Interactive Video Network: An Overview of the Video Manager and the V Protocol

Roger E. Libman
Mowaffak T. Midani
Irene J. Morgan
Hung Tan Nguyen

The AT&T Interactive Video Network (IVN) offering consists of the network elements and signaling protocol necessary to provide interactive services such as video on demand, home shopping, and interactive games, while establishing a framework for evolution to support future services. This paper introduces the network architecture being developed and the two most frequently deployed access technologies—switched digital video (SDV) and hybrid fiber-coax (HFC). It then focuses on two major AT&T Network Systems offerings that support the interactive video network—the AT&T Video Manager and the V Protocol.

Introduction

Movies on demand, video conferencing, distance learning, and telemedicine are just some examples of interactive visual and multimedia applications that are rapidly becoming realities. They promise to enhance the quality of communication in the years to come. While establishing an interactive video network (IVN) presents many challenges, the availability of enabling technologies and standards such as the synchronous optical network (SONET), the asynchronous transfer mode (ATM), and Motion Picture Experts Group (MPEG)-2 (see Reference 1) are helping to bring it closer to reality. (See Panel 1 for definitions of abbreviations, acronyms, and terms.)

The architecture AT&T is deploying provides an infrastructure for delivering interactive multimedia services, one of which may be video. Based on a broadband switched ATM network, this network architecture uses two access technologies[2] most frequently: hybrid fiber-coax (HFC) and switched digital video (SDV). With either access technology, this network can provide the entire range of interactive services. Two key elements of this architecture are the AT&T Video Manager and the user-to-network protocol—the V Protocol—used by the Video Manager to communicate at the session layer with video information provider's (VIP's) servers and end users.

The Video Manager, a key network element in an end-to-end Video Dialtone platform, is the central intelligence and control for the interactive video network, performing network resource allocation and management, session and connection management, and measurements collection. Its Level 1 gateway functionality satisfies the Federal Communications Commission's (FCC's) Video Dialtone requirement for equal access to any video information provider. It is the end user's initial access to interactive services.

The V Protocol is a user-to-network signaling protocol at the session layer that supports advanced interactive applications. It was introduced by AT&T with the following philosophy in mind: network technology independence above the transport layer; clear distinction between the concepts of "sessions" and "resources"; abstract representation of network resources; and natural co-existence with the Q.931/Q.2931 signaling protocol suite. To make this protocol an open industry standard for interactive services, AT&T, in cooperation with other parties representing the telecommunications, cable, and computer industries, has made international standards submissions based on the V Protocol to the MPEG standards body.

This paper presents a high-level view of the network architecture and the two access technologies, and then describes the Video Manager and the V Protocol in more detail.

Purpose of the Interactive Video Network. The IVN is designed to provide a variety of interactive services. Variations of the basic network architecture are being used by local exchange carriers (LECs) and cable television (CATV) operators. Typical interactive services include movies on demand, interactive games, home shopping, bank at home services, and cable services that offer broadcast channels, premium channels, and pay-per-view. The network can also provide telephony services, including "plain old telephone service" (POTS), integrated services digital network (ISDN), and video telephony. This paper emphasizes interactive video services.

Network Architecture. Figure 1 depicts a high-level view of the network architecture. It shows geographically dispersed *media servers*; a wide area broadband-switched ATM network; multiple access nodes, each with distribution networks; and residential equipment such as set-top terminals (STTs) or personal computers. Figure 1 also shows two alternate access technologies for the distribution portion of the network—a trunk and branch HFC structure similar to a traditional CATV network, and a star network, also referred to as switched digital video (SDV). These are described briefly below.

The Video Manager has centralized control of the network. In the LEC environment, the FCC has mandated that the network maintain an open interface to allow multiple nonregulated VIPs to attach to the network. End users must have equal access to these VIPs. A function called the Level 1 gateway enables users to choose and connect to a desired VIP in real time. The Level 1 gateway capability is one key role of the AT&T Video Manager.

The network contains many elements, including:
- The media server,
- The broadband switched ATM network,
- The Video Manager, and
- The access node and distribution.

The media server. Though the architectures of media servers vary, all such servers maintain a library of program material (such as movies) stored on disk and/or magnetic tape. Users are able to watch different programs or different parts of the same program simultaneously and to exercise "VCR-like" control capabilities, such as pause and fast forward. Programs are typically stored in a compressed digital format, such as MPEG-2 transport streams, at compression rates of 1.5 to 6 Mb/s. When a program is played out of a media server, the MPEG packets are encapsulated in ATM cells and transported over the switched ATM network. By allowing applications to select video segments and data downloads based on user choices, more complex applications, such as home shopping and interactive games, can also be provided.

The broadband-switched ATM network. Media servers are connected to access nodes through the broadband switched ATM network. The network is generic, not video-specific, and is not "aware" that video is being transported in the ATM cells. The switched ATM network may, in fact, be shared by other unrelated services. The

Panel 1. Abbreviations, Acronyms, and Terms

ATM—asynchronous transfer mode
CATV—cable television
DSM-CC—Digital Storage Media Command and Control
FCC—Federal Communications Commission
GUI—graphical user interface
HDT—host digital terminal
HFC—hybrid fiber-coax
ISDN—integrated services digital network
ISO—International Organization for Standardization
IVN—interactive video network
LEC—local exchange carrier
MPEG-2— Motion Picture Experts Group-2
NTSC—National Television Systems Committee
OA&M—operations, administration, and maintenance
OSS—operations support system
PIN—personal identification number
POTS—"plain old telephone service"
PVC—permanent virtual circuit
QoS—quality of service
RF—radio frequency
SONET—synchronous optical network
SDV—switched digital video
STT—set-top terminal
SVC—switched virtual circuit
UNI—user-network interface
VIP—video information provider

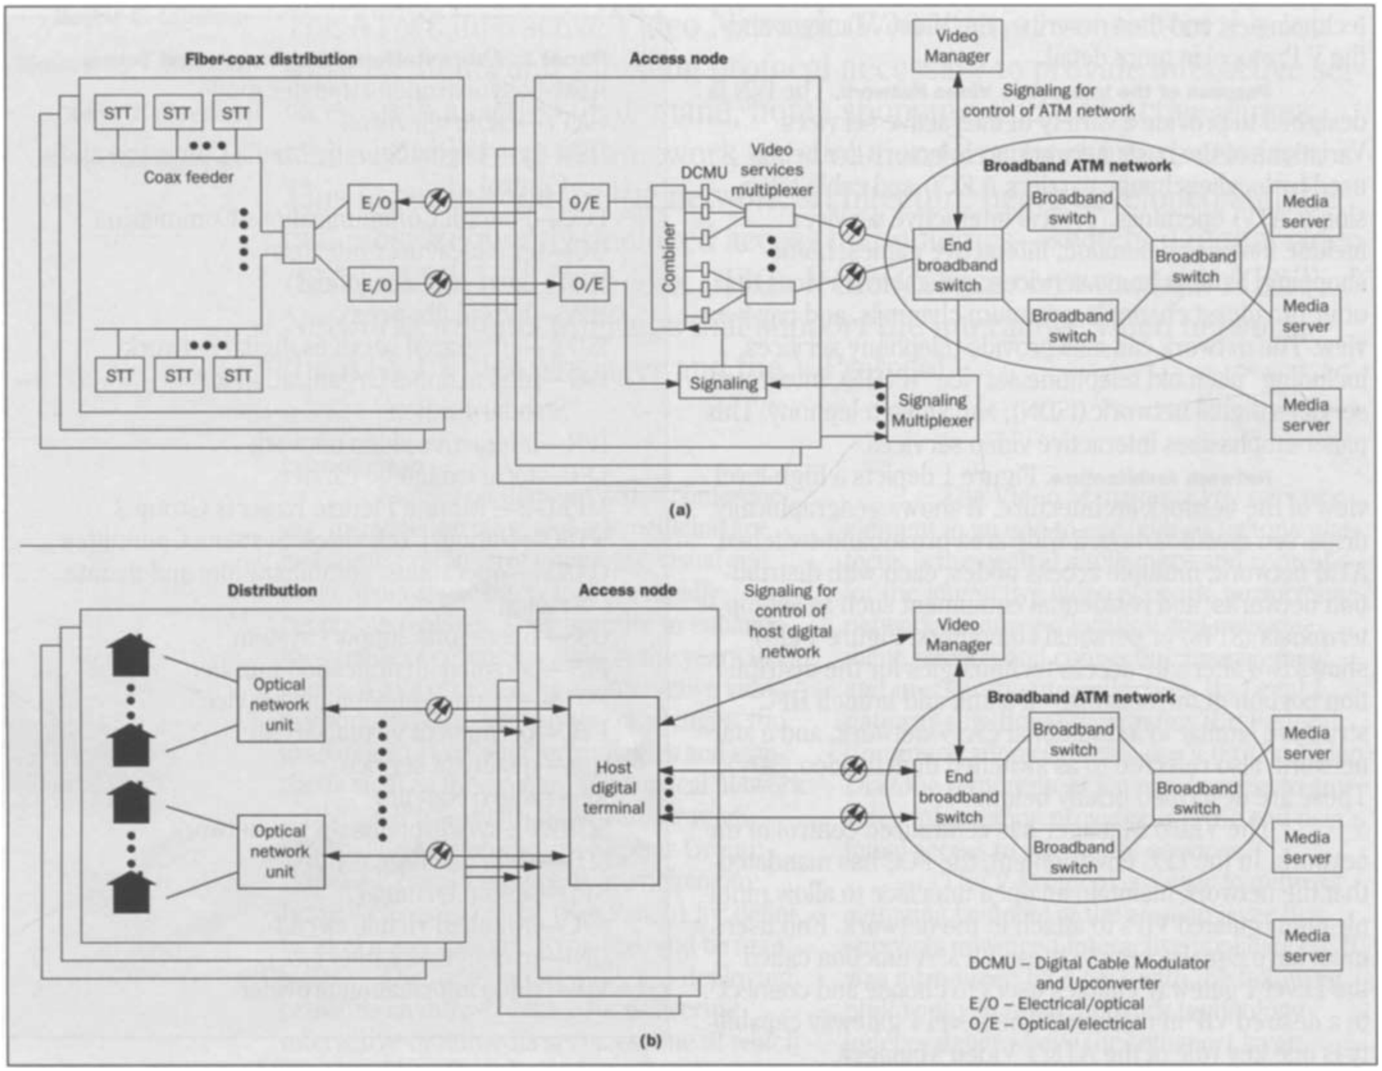

Figure 1. High-level network architecture of the IVN, including (a) the HFC access network (trunk and branch), and (b) the SDV access network (fiber to the curb).

AT&T switched ATM product is a member of the GlobeView™-2000 broadband system offering. The present version offers a 20-Gb/s fabric, switched virtual circuits (SVCs), and permanent virtual circuits (PVCs).

The Video Manager. The Video Manager component, described later in more detail, plays a number of roles. The first is connection and session management. Another is the Level 1 gateway mentioned earlier. On request of the VIP or end user, the Video Manager establishes ATM connections for application download, video from the server to the STT, and end user control functions. To accomplish this, the Video Manager provides the signaling capability to control the switched ATM network, establishing ATM connections on behalf of end users and VIPs. The Video Manager maintains records of all ATM connections used within an interactive session and compiles them as needed to bill the VIP for the use of the network.

Access node and distribution. Two architectures are used for the access node: the HFC (trunk-branch topology) and the SDV (star topology).

The HFC network topology is similar to that of conventional CATV systems. Radio frequency (RF) subcarrier modulation is used to transport signals through the access network. In the United States, these are transported in 6-Mhz channels in National Television Systems Committee (NTSC) format. Other intervals (for example, 8 Mhz) can also be used, depending on the television standards. A tuner in the STT selects the appropriate channel. The major difference between this and conventional CATV networks is the addition of compressed digital video. In the CATV network, one channel of analog video occupies 6 MHz. With digital compression, multiple digital programs can be placed in one channel, thus significantly increasing the capacity of the system. Such a system can carry hundreds of digital programs, allowing for point-to-point interactive services such as movies on demand. Because the transmission medium in this access architecture is shared, encryption is used to control access and privacy. Another difference between this and conventional CATV networks is more sophisticated upstream signaling.

The components of the access network are responsible for:

- Multiplexing of digital programs that occupy a 6-MHz channel,
- Modulation and combining of the channels into the RF spectrum for transport on the fiber-coaxial distribution network,
- Routing of signaling messages to-from the STTs, and
- Electrical-to-optical conversion for transport over the optical portion of the distribution network.

The second type of access network architecture is SDV, in which digital information is transmitted directly as "bits" through a dedicated path to the end user. The access network is configured as a star topology, which is similar to a telephony fiber-to-the-curb architecture, with modifications to support digital video services. Unlike the HFC trunk and branch architecture described earlier, this approach provides a separate physical drop, or line, for each home, with a dedicated downstream bandwidth on the order of 45 Mb/s.

In the HFC network, video content for a specific home in a neighborhood is present on the coax feeding all homes. Each STT is responsible for tuning to the appropriate 6-MHz channel and decoding a specific program, while ignoring all other information. In contrast, in the SDV network, as in a conventional telephone system, only the content required by a home is actually delivered to it. The available 45-Mb/s bandwidth allows for multiple STTs in a home, each receiving a unique program. It is not possible, however, to have a large number of programs (a full complement of broadcast and cable channels) entering the home at the same time. Switching digital broadcast channels is done outside the home. A channel change request originating at the STT sends an upstream message to the ***host digital terminal (HDT)***, where all broadcast channels are available. Similarly, for point-to-point services like movies on demand, switching must take place in both the HDT and the broadband ATM network.

The HDT is responsible for:

- Initializing/booting each STT as it is powered up,
- Handling upstream/downstream signaling for all interactive services,
- *For broadcast services*, receiving provisioning information from VIPs that determines which program streams an STT is authorized to decode,
- Handling end user requests for broadcast service by receiving a message from the STT and connecting the STT to the appropriate program stream,
- Routing end user requests for point-to-point interactive services to the Video Manager, and
- Acting on commands from the Video Manager to connect an STT to the broadband ATM network for point-to-point interactive services.

In addition, the HDT includes SLC® carrier hardware, used to connect end users to the narrowband telephony network.

Comments on the Support of Both Topologies

Each of these access topologies has its role in the marketplace, depending on the needs of the service provider. Many factors influence the choice, such as outside plant environment, living unit density, installed base of equipment, expected service penetration rates and growth rates, and expected operations methods and costs. The optimum choice depends on all these factors; therefore, AT&T supports both topologies. The signaling philosophy described next makes doing so practical.

Network Signaling Philosophy. The network signaling scheme is designed to isolate the STT from the details of the network. This has two major advantages:

First, it allows the STT to work with networks of varying complexity without the need to be customized for each. The network may be HFC or SDV. It may by based on ATM SVCs (as described here), PVCs, or some technology other than ATM, without the STT having this knowledge. A network could be migrated from one technology to another without change to the STTs. Second, it relieves the STT of handling complicated protocols, thus reducing memory and processing requirements and their associated costs. Because the STT is the most replicated entity in the network, even a small cost savings per STT can mean significant cost reduction for the overall network.

Session Layer Protocol. The STT and server are isolated from the details of the network control signaling by introducing a higher-layer protocol through which the STT and server communicate with the Video Manager. This is a *session layer* protocol, as opposed to a *network layer* protocol, which directly controls the switched network.

In this context, the term *session* has this specific meaning: In interactive applications, it is often necessary to create several connections through the network between server(s) and a particular STT. For a movie-on-demand application, the connection over which the end user interactively chooses a movie to view may be bidirectional, but the content may be transported on a downstream high-bandwidth connection. The properties of connections may vary in such aspects as directionality (unidirectional, bidirectional), symmetry/asymmetry, bandwidth, and quality of service (QoS). A session implies a relationship between an end user and a VIP for a period of time during which the end user participates in interactive activity. Associated with that session are the series of related connections through the network between server(s) and the user's STT.

To allow the STTs and servers to control the network at the session level, the Video Manager assumes the roles of session manager and proxy signaling agent. STTs and servers interact with the Video Manager through session-level protocol to establish sessions and add resources (such as connections) to a session. The Video Manager provides the network-level signaling, as necessary, depending on the design of the specific network, and acknowledges success or failure to the STT and server.

The Video Manager

The AT&T Video Manager is a key network element of an end-to-end network architecture. As the central intelligence and control for the IVN, it provides an end user's initial access to interactive services. It is the trusted point of control for network resource allocation, connection management, and measurements collection necessary for session management and billing data services. It guides the viewer's service selection and access, allowing the viewer to choose from the services available.

The Video Manager provides the Level 1 gateway functionality necessary to satisfy the FCC's Video Dialtone requirement[3] for equal access to any VIPs, using open interfaces and standard protocols, where available. Where standards do not exist, AT&T is making submissions to the appropriate organizations to support standardization efforts.

Video Manager Software Architecture. The software architecture of the Video Manager is modular, allowing multiple network architectures to be built on a common software platform. This is critical in an emerging market that has not converged on a common architecture. The Video Manager supports both SDV and HFC network architectures using the same software platform. Its functionality can also be distributed if additional performance management or load isolation is desired.

Video Manager software subsystems use a standard UNIX* platform and a set of libraries to provide software independence from the platform hardware. The Video Manager software has four major subsystems: communication protocols (specific to STTs, servers, and network hardware); session and connection management; subscriber services; and billing, operations, and network management interfaces.

Each network element that the Video Manager communicates with has individual modules in the communications protocols subsystem. This provides a firewall for the rest of the Video Manager, while allowing changes and additions to the network. The modules that comprise the session and connection management subsystem perform the session, connection, and resource allocation functions. The subscriber services subsystem supports applications such as the menu of services. The billing, operations, and network management interfaces subsystem provides system administration functions and communicates with centralized operations support systems (OSSs).

The Video Manager software foundation and a commercial UNIX platform support these subsystems. The Video Manager's software foundation provides encapsulation and common support services that allow its application layer software to be hardware independent and distributed across multiple platforms.

Video Manager Functionality. The Video Manager provides:

- Network communications,
- Session and network resource management,
- Subscriber services, and
- Network provider support.

The network communications function provides the protocols to communicate with other network elements.

The session management function includes session set-up and tear-down. It also establishes and maintains session records for all active sessions and provides usage-related counts to the OSS infrastructure for network engineering and billing purposes. The network resource management function performs resource allocation for network resources such as bandwidth. The Video Manager also recovers network resources after network element outages.

Using its subscriber services capability, the Video Manager is the end user's initial entry into the network and guides the end user in service selection. This provides the Level 1 gateway role mandated by the FCC for LEC networks. The Video Manager also maintains a profile database, including preselected services information and any viewing restrictions set by parents.

The Video Manager provides essential services to network providers, with tools for service provisioning, network access and billing data collection, and status monitoring. Because the Video Manager stores information related to end users, it can serve as a central repository for end user data and can provide this information to VIPs and other network elements, creating revenue opportunities for network providers.

Network communications. The Video Manager communicates with other network elements using the appropriate session layer protocol, as determined by the network provider. The V Protocol, described in this paper, provides a feature-rich capability with the necessary functionality to support session layer communication.

Session management. The Video Manager performs the functions of session management, network resource management and allocation, and bandwidth management for the network. It correlates and tracks the resources used by a session across the network, including session establishment, session ongoing services, normal session release, abnormal session termination, and session auditing. During a session, it can add and drop resources, and even cancel the session.

A session may need to be cancelled if communication is lost between the server and the STT, if no upstream signaling path is established, or if the end user directly requests that a session be cancelled.

At the ATM network layer, the Video Manager provides sessions with connection management services. To set up the various paths for the session, the Video Manager, STT, and network access equipment use the network-level signaling protocol (Q.2931) to establish a session-based connection between the Video Manager and the STT. Control messages are exchanged between the STT and the Video Manager along this connection, enabling the Video Manager to determine the desired service for the particular session. The Video Manager informs the selected VIP of the request made by the STT and sets up a network connection between the STT and the VIP. It works with the network access elements—using network layer signaling protocols—to continue with the network connection setup through the access network. This sets up a control path between the STT and the server. The data path is established in a similar manner.

The Video Manager has an authentication feature that detects whether the hardware address of the residential customer premises equipment matches its network point of presence. The Video Manager will block service if these values do not match. This feature can be turned on or off. If the parameter is set to block service on an authentication failure, the appropriate actions for denying service, as defined by the communication protocol, are followed.

Network resource management. In response to requests to add resources, the Video Manager allocates resources to a session. It maintains a table of available resources and makes assignments of the available coax subchannels in an HFC network; it also provides mapping between the assigned subchannel and the ATM PVC or SVC allocated to the session. By doing so, data and video paths are established between the end user's STT and the selected VIP. The Video Manager sets up the signaling path between the STT and the server to allow user-to-server communication. It also allocates other network resources, including access encryption keys, which control access and privacy on the HFC network.

Subscriber services. For digital services, the Video Manager offers a service directory (also called a Level 1 menu). During the session, an STT may receive the menu information from the Video Manager, which in turn can deliver one of several different service directories based on the service model and the needs of the end user. Service models include:

- *Quick-connect.* This automatically connects the end user to a designated service provider.
- *Service directory.* This service model sends the end user a service directory.
- *Identify end user first.* This model requires the end user to self-identify by entering a personal identification number (PIN) before receiving service. With this service model, the Video Manager can support multiple end users for a single STT-id. Each PIN for a specific STT-id can have different network treatments, such as preferred VIP or VIP blocking.

A service directory is distributed using the session-based service control channel or, for larger applications, a dedicated channel. This directory information is presented by the STT, which may, as an application design issue, combine it with graphics, animation, or audio effects.

After the end user makes a selection, the STT sends the selection identifier to the Video Manager. When the Video Manager receives the selection identifier, it contacts the selected VIP and requests a session for the end user using the appropriate session layer protocol.

The Video Manager supports blocking enforcement, which enables each end user to choose between the general menu and a restricted menu for his or her default services directory. The restricted menu is data driven and can be used to provide end users with a directory option that does not show potentially objectional offerings (such as adult entertainment).

Network Provider Support. Measurements provided by the Video Manager include session and resource data for billing purposes, as well as network and service usage information. To bill VIPs for their network access and usage, records for all resources allocated for a session are stored by the Video Manager for network access billing use. The data can be formatted on a remote billing system and retrieved on demand or at a scheduled time, or it can be processed locally. Billing records contain information pertinent to network usage billing, such as a session identifier, a structure code indicating service completion status, originating STT address, selected VIP, billing party, resources used, bandwidth allocated, type of service, time and date service started, and duration.

The Video Manager collects and stores network traffic usage information, which is used in engineering and administration of the Video Manager and the network. Traffic measurement data is available to network administrators and system engineers in several formats: on-demand or scheduled reports and batch data files, and on-demand display screens.

Data collected on the network and its components includes information such as known outages and capacity utilized over time.

The Video Manager provides information about itself and its status to a centralized surveillance system. Simple network management protocol agents are used for alarms and notification. The Video Manager also maintains its own view of the status of network resources, such as VIPs, STTs, communication links, and channels.

Several interface choices for operations, administration, and maintenance (OA&M) are provided by the Video Manager. Local interfaces include a command line interface, as well as a graphical user interface (GUI). All OA&M services can be accessed remotely using standard remote login and X-Window System* services, which allow remote operations from a centralized surveillance facility. The Video Manager GUI is an X-Windows-based interface with a Motif* "look and feel." It provides network operations staff with an easy to use, consistent interface to key Video Manager data and functions. Windows for Video Manager maintenance and service administration, such as billing and measurements administration, are also included in the Video Manager GUI.

Provisioning, or bulk loading, of network resource and end user data onto the Video Manager's database can be performed using the OSS link, or locally using text commands or a graphical forms interface. Primary Video Manager provisioning is conducted through an external communications interface to an OSS, such as the AT&T Video Services Manager. The data pro-

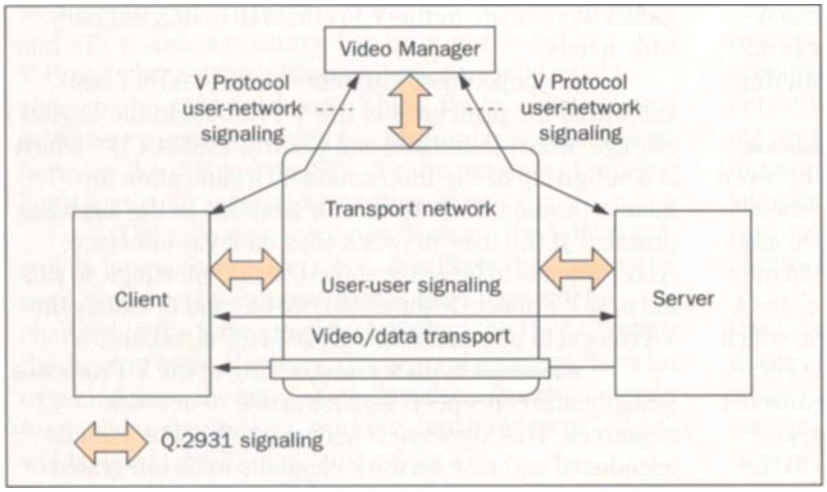

Figure 2. Signaling network reference model showing V Protocol and Q.2931 message flows.

visioned typically includes information on end users, STTs, and VIPs. The Video Manager also contains data describing all downstream network elements and preassigned signaling paths.

The V Protocol

The V Protocol is a message-based signaling protocol that is used to establish, manage, and tear down sessions between VIPs (servers) and end users (clients) over the IVN.[4] The V Protocol defines the messages, scenarios, and resource structures for supporting session-level management functions, such as:

- Establishing and terminating sessions, as requested by the client or server;
- Dynamic resource allocation/deallocation during the lifetime of a session, as requested by the client or server;
- Forwarding and transfer of sessions and resources between servers and/or applications, including support of a distributed application environment; and
- Status and other OA&M activities for monitoring the state of a session.

Signaling Reference Model. The V Protocol introduces a session layer set of messages and procedures that describes the process of establishing and releasing sessions/connections across the network. For networks that support switched services, the V Protocol provides a mechanism that uses the Q.2931 messages at the network layer to establish specific virtual circuits. Figure 2 shows a signaling reference model for the IVN.

As shown in Figure 2, clients and servers establish sessions by interacting with the Video Manager over a transparent transport network. For switched services, Q.2931 messages are used to set up and release switched virtual circuits, under the direction of the V Protocol session control.

After a session is established, direct signaling between the client and server occurs via the user-to-user signaling protocol, without involving the Video Manager. Video and data are transferred over a unidirectional or bidirectional pipe from the server to the client. User-to-user signaling protocol is beyond the scope of this paper.

V Protocol Philosophy. AT&T has developed the V Protocol with the following philosophy in mind:

- Network technology independence,
- Distinction between the concepts of *sessions* and *resources,*
- Distinction between the concepts of *calls* and *sessions,*
- Abstract representation of resources,
- Dynamic resource allocation,
- Complementary role with other control protocol, and
- An open industry standard.

Network technology independence. The V Protocol is not linked to any particular transport network technology. It uses any available network layer to establish and manage communication paths between endpoints.

Distinction between concepts of sessions and resources. The V Protocol makes a clear distinction between sessions and network resources. While a session is an association between parties (a server and a client), a resource is a trackable "object" or "element" allocated and retrieved by the Video Manager within the session. This distinction allows for the development and offering of advanced applications to the end user.

Distinction between the concepts of calls and sessions. The ATM Forum[5] defines a call as an association between end users or between an end user and the ATM network. In the context of the V Protocol, a session is also an association, but in broader terms. While a call is defined only in the context of an ATM network, a session associates a server and a client in the context of a IVN network, which may include an ATM network as a subcomponent. Therefore, the session association may include ATM transport resources, other (non-ATM) transport resources, and other types of resources that may not be related to the transport capability. A session may not need any one-to-one connection at all, such as a "broadcast" session or a "datagram-type" session. For the IVN network provider, a session is used as the point of entry to implement higher-level policy decisions and enhanced IVN services (such as "Session Transfer"), and to monitor all resources related to a current session between a client and a server.

Abstract representation of resources. The V Protocol abstracts network resources and views them as objects with unique characteristics. By abstracting resources, new network architectures can be introduced and new network elements can be integrated into existing networks without the need for major changes to the protocol.

Dynamic resource allocation. The V Protocol supports dynamic resource allocation during the lifetime of a session. Consequently, resources can be added to or deleted from an active session according to the needs of the client, server, and the application being run.

Complementary role with other control protocol. A main objective of the V Protocol philosophy was to avoid competing with or replacing existing network-level protocols, such as Q.2931. At the session layer in the protocol stack, the V Protocol supports resource negotiation and advanced session features, while the Q.2931 network protocol establishes calls and switched connections. In that regard, the V Protocol and the Q.2931 protocol play complementary roles in supporting advanced interactive services.

Open industry standard. Designed to be an open industry standard, the V Protocol was made available to many parties in the telecommunication and cable TV industry. Based on the feedback received, enhancements were made to the V Protocol to match industry-wide needs.

In cooperation with other parties, AT&T has submitted the principles of the V Protocol to the Digital Storage Media Command and Control (DSM-CC),[6-9] which is a subgroup of the International Organization for Standardization (ISO) MPEG, for adoption as the signaling protocol at the user-network session layer interface. AT&T intends to be active at the DSM-CC meetings, to submit new V Protocol features to DSM-CC, and to modify the V Protocol to be consistent with DSM-CC directions.

Resources in the V Protocol. One of the V Protocol's fundamental concepts is its abstraction of network resources. This allows new network architectures to be introduced and new network elements to be integrated into existing networks without the need for major changes to the protocol. Resources are requested in the form of *resource descriptors* and are allocated to the VIPs, STTs (clients), and other transport and signaling entities, also in the form of resource descriptors.

For HFC and SDV networks, the V Protocol supports resource negotiation between the user (VIP or STT) and the Video Manager. Currently, most network architectures only allow the VIP to request and negotiate resources with the Video Manager; however, the V Protocol can support a symmetric network in which the VIP or the STT can request and negotiate resources with the Video Manager.

An HFC system uses RF subcarrier modulation to transport signals through a shared access transport network. An SDV system transmits the digital voice, video, and data signals as bits directly through a dedicated transmission path to the end user. The two transport networks have different resources and thus require different resource descriptors. However, the general structure of the resource descriptors is universal.

V Protocol Scenarios. The V Protocol's sequence of commands can be described in terms of scenarios. The basic scenarios related to session management include *session request* and *session release* scenarios. The basic scenarios related to resource management include *resource request* and *resource deletion* scenarios.

The International Telecommunications Union (ITU) defines Q.2931 as the signaling protocol to set up ATM connections. This protocol is used with the V Protocol if the requested resource is an ATM connection. The Q.2931 protocol sets up a switched virtual circuit between the VIP and STT if such a resource has been requested in a V Protocol message. Other protocols may be used to allocate non-ATM resources. The V Protocol sets up and maintains a session, which is a higher-layer association between the VIP and the STT, for the purpose of receiving a particular interactive multimedia service.

The following scenarios illustrate the V Protocol and its interactions with Q.2931. It is likely that the message names and parameter names used here will be changed in the future to reflect work of the DSM-CC standardization body. However, the general message flows are expected to be the same. To follow the current nomenclature used in the DSM-CC working draft,[6] the term "client" will be used for STT, and "server" for VIP.

Session Setup. When an end user selects a multimedia service, the client sends a Session Request message to the Video Manager. This message contains a *terminal-id,* which identifies the client, and a *service-id,* which identifies the selected interactive service. (See Figure 3 for a description of a client session setup.) On receipt of this message, the Video Manager makes various high-level decisions, such as rejecting requests from clients of nonpaying end users or to servers of nonpaying VIPs. These decisions are based on the policies set up by the network provider. After the Video Manager has validated both the client and server, it contacts the indicated server using the low-level transport and network addresses of the server (e.g., IP address or E.164 address) that the Video Manager maintains. Tthe client only knows about the high-level service-id.

The server receives the Session Indication message from the Video Manager and makes the decision to accept or reject the request. If the server accepts the request, it responds to the Video Manager with a Session Response message, indicating any resources that it needs from the network to support the selected service. If the service needs ATM SVC connections, the server will include the proper number of "ATM SVC" resource descriptors. These resource descriptors contain sufficient information for the Video Manager to later establish a nonassociated connection between the server and the client or some other downstream signaling node that terminates ATM.

The Video Manager receives the Session Response message and processes the resource descriptors included in the message. For an "ATM SVC" resource, the Video Manager verifies the requested properties, such as user cell rate and QoS, against the network provider's high-layer policies, such as rejecting requests of bandwidth greater than 10 Mb/s during evening hours. The Video Manager is also responsible for translating parameters in the "ATM SVC" resource descriptors into Q.2931 information element parameters.

For each valid "ATM SVC" resource request, the Video Manager constructs a replying resource descriptor, collects all assigned resource descriptors in a Resource Indication message, and sends it to the server. The server acknowledges with a Resource Response message.

If the *response* field in the server Resource Response message is positive for each "ATM SVC" resource requested, the Video Manager will initiate a Q.2931 nonassociated Call/Connection Procedure.

The Video Manager initiates a nonassociated Call/Connection Procedure by sending a Q.2931 Setup message to its ATM user-network interface (UNI). The Call Reference used will be selected by the Video Manager. The calling party number is the ATM address of the client as derived from the terminal-id field; the called party number is the ATM address of the server as derived from the service-id field or received in the "ATM SVC" resource descriptor of the Session Response message. This feature allows a server to redirect the recipient of the ATM call to any distributed endpoints it manages. The ATM adaptation layer parameters, ATM user cell rate, and QoS parameter contain values derived from the "ATM SVC" resource request descriptor.

When the server is informed of the SETUP on its ATM UNI, it will associate the Call Reference with this session layer's information. The server maintains this association until the ATM SVC connection is released. The Video Manager maintains the relationship between the Call Reference at its UNI and the session-id/resource-id pair to maintain its ability to retrieve either one from the other.

After all requested resources have been assigned, including the "ATM SVC" resources, the Video Manager informs the client through the V Protocol session layer's client Session Confirmation

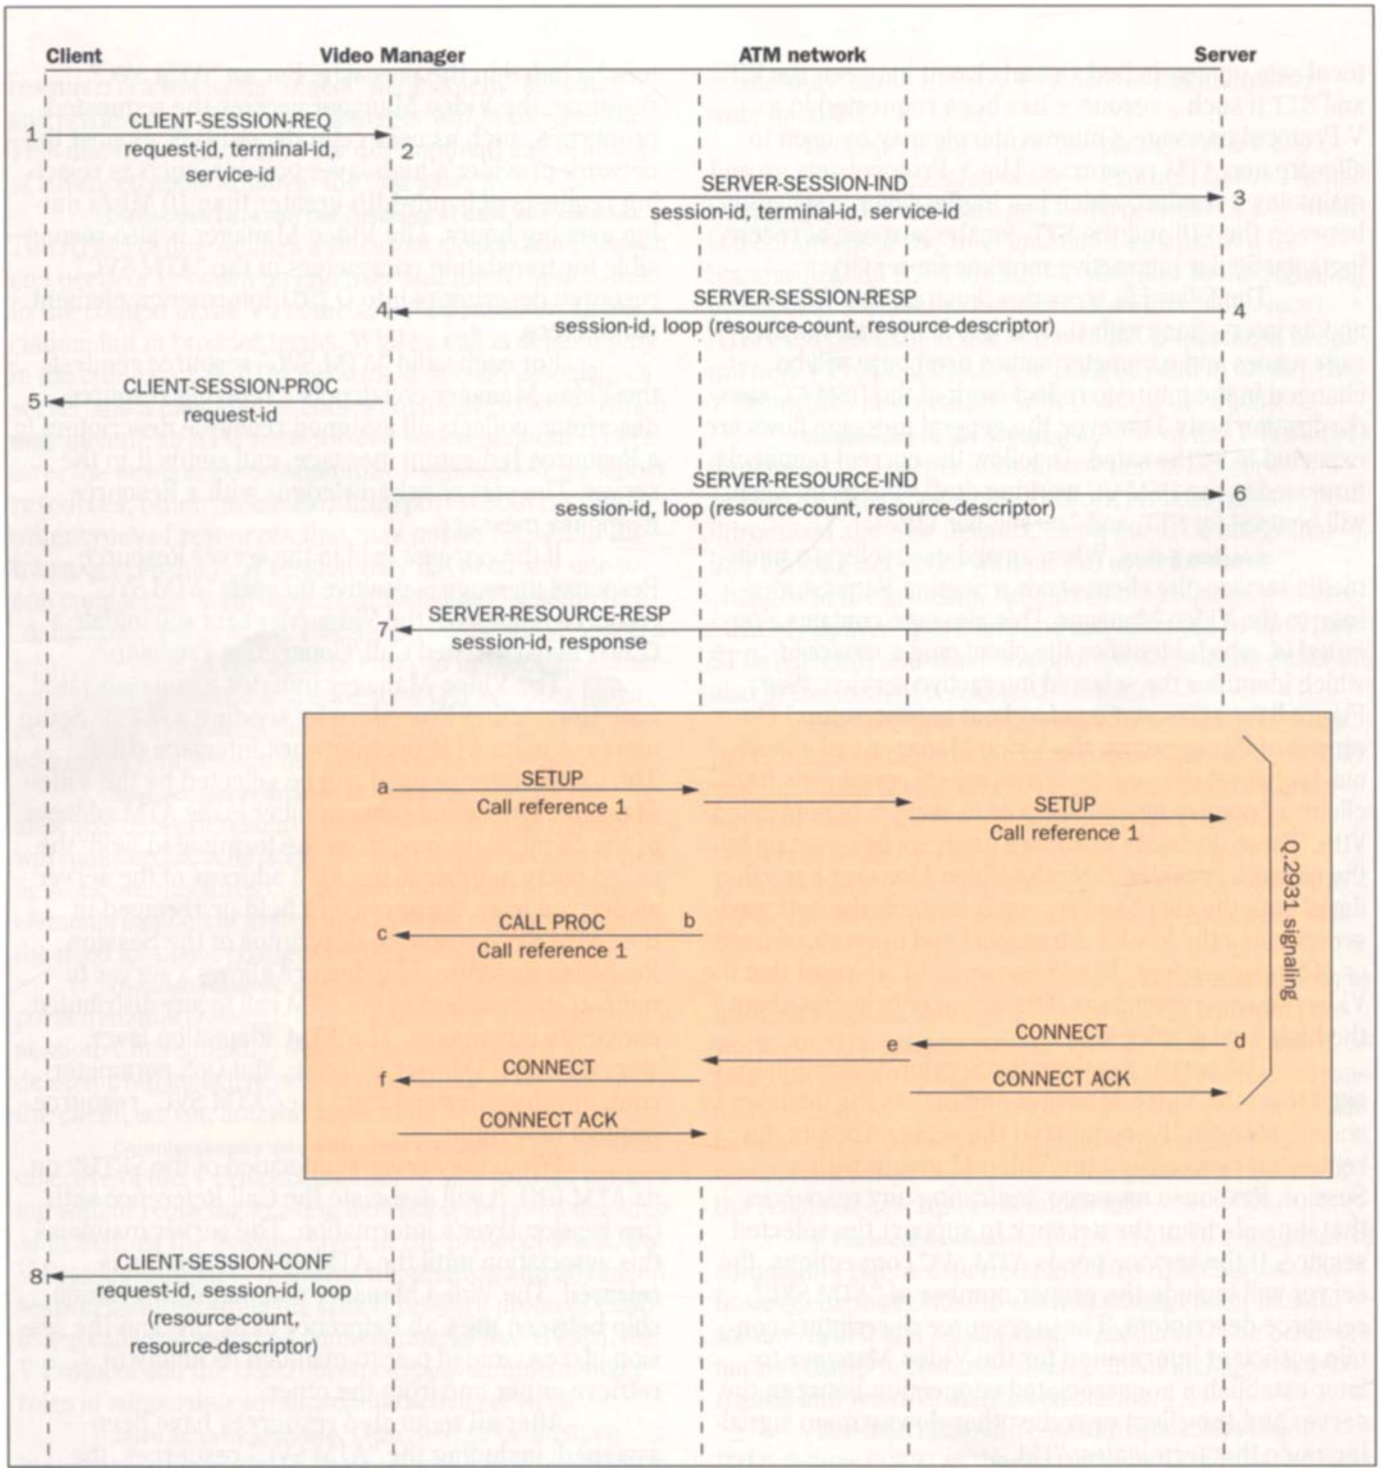

Figure 3. Client-initiated session request scenario.

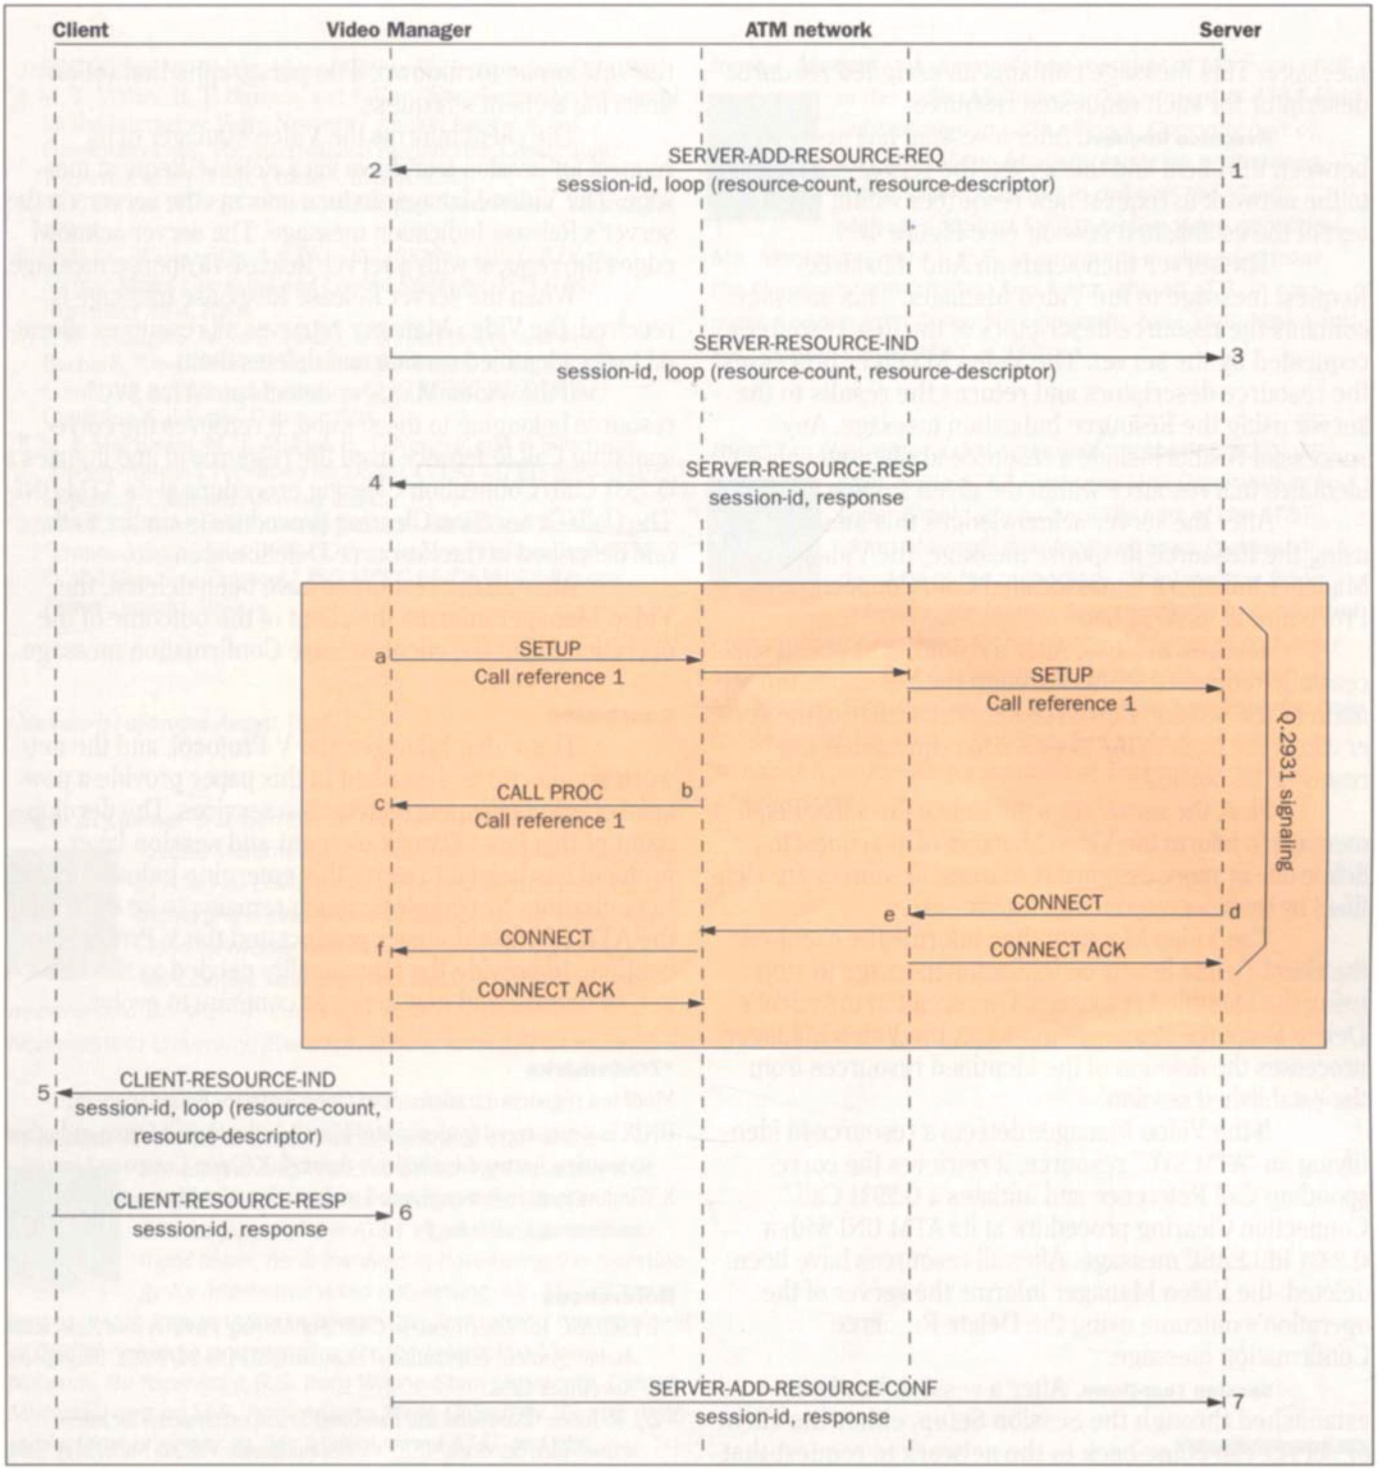

Figure 4. Server resource request scenario.

message. This message contains an assigned resource descriptor for each requested resource.

Resource Request. After a session has been set up between the client and the server, the server can return to the network to request new resources within the context of the established session (see Figure 4).

The server then sends an Add Resource Request message to the Video Manager. This message contains the resource descriptors of the new resources requested by the server. The Video Manager processes the resource descriptors and returns the results to the server using the Resource Indication message. Any successful results include a resource-id that uniquely identifies that resource within the given session.

After the server acknowledges this message using the Resource Response message, the Video Manager initiates a nonassociated Call/Connection Procedure at its ATM UNI.

Resource Deletion. After a resource has been successfully requested, either through the Session Setup scenario or the Session Resource Request scenario, the server can come back to the network to request that the resource be deleted.

First, the server uses the Delete Resource Request message to inform the Video Manager of its request to delete one or more assigned resources. Resources are identified by their corresponding resource-id.

The Video Manager then informs the client via the client Delete Resource Indication message to stop using the identified resources. On receipt of the client's Delete Resource Response message, the Video Manager processes the deletion of the identified resources from the established session.

If the Video Manager detects a resource-id identifying an "ATM SVC" resource, it retrieves the corresponding Call Reference and initiates a Q.2931 Call/Connection Clearing procedure at its ATM UNI with a Q.2931 RELEASE message. After all resources have been deleted, the Video Manager informs the server of the operation's outcome using the Delete Resource Confirmation message.

Session Tear-Down. After a session has been established through the Session Setup, either the client or server can come back to the network to request that the session be torn down. The paragraphs that follow describe a client's request.

The client informs the Video Manager of its request for session tear-down via a Release Request message. The Video Manager in turn informs the server via the server's Release Indication message. The server acknowledges the request with a server Release Response message.

When the server Release Response message is received, the Video Manager retrieves all resources allocated to the identified session and deletes them.

If the Video Manager detects an "ATM SVC" resource belonging to the session, it retrieves the corresponding Call Reference from the resource-id and initiates a Q.2931 Call/Connection Clearing procedure at its ATM UNI. The Call/Connection Clearing procedure is similar to the one described in the Resource Deletion scenario.

After all the resources have been deleted, the Video Manager informs the client of the outcome of the operation using the client Release Confirmation message.

Conclusion

The Video Manager, the V Protocol, and the network architectures described in this paper provide a powerful offering to support interactive services. The development of this key network element and session layer protocol has helped to move this emerging industry closer to realization. Nonetheless, much remains to be done, and the AT&T Video Manager product and the V Protocol will continue to provide the functionality needed as the market, standards, and new services continue to evolve.

*Trademarks

Motif is a registered trademark of Open Software Foundation, Inc.

UNIX is a registered trademark of Novell in the United States and other countries, licensed exclusively through X/Open Company Limited.

X Window System is a registered trademark of the Massachusetts Institute of Technology

References

1. *ISO/IEC 1381801:Generic Coding of Moving Pictures and Associated Audio: Systems*, International Standard, MPEG 94/N801, Singapore, November 1994.
2. J. R. Jones, "Baseband and Passband Transport Systems for Interactive Video Services," *IEEE Communications*, Vol. 32, No. 5, May 1994.

3. FCC Docket No. 91-334, *Video Dialtone Order*, November 27, 1991.
4. M. T. Midani, H. T. Nguyen, and S. Ger, "Introducing the V Protocol for the Interactive Video Network," *Annual Review of Communications*, The International Engineering Consortium, Western Cable Forum, Chicago, Illinois, 1995.
5. ATM Forum, *ATM User-Network Interface Specification*, Version 3.0, 1993.
6. DSM-CC, *Working Draft of ISO/IEC 13818-6: MPEG-2 Digital Storage Media Command and Control Extension*, MPEG 94, November 1994, N806.
7. T. H. Addington, M. S. Goldman, M. J. McPheters, and K. E. Bechard, "User-Network Primitives for Digital Storage Media Command and Control Extension," *ISO MPEG 94/DSM-01*, Cupertino, California, October 1994.
8. M. J. McPheters, M. T. Midani, H. T. Nguyen, and P. Silverman, "User-to-Network Primitives for DSM-CC, ISO MPEG 94/DSM-010," Cupertino, California, October 1994.
9. AT&T, US West, Southwestern Bell, and Broadband Technologies, "Proposal to Combine DSM-CC Session Layer Protocol and ATM Q.2931 Signaling Protocol," *ISO MPEG 95/DSM-022*, Rennes, France, January 1995.

(Manuscript approved August 1995)

Roger E. Libman *is a member of technical staff in the Visual/Multimedia Communications business unit at AT&T in Lisle, Illinois. He is currently working on defining the end-to-end architecture for interactive video networks with an emphasis on signaling. Mr. Libman, who received a B.S. in electrical engineering and an M.S. in computer science, both from Northwestern University, Evanston, Illinois, joined AT&T in 1969.*

Mowaffak T. Midani *is a member of technical staff in the Broadband and Photonic Switching Department at AT&T Bell Laboratories in Naperville, Illinois. Currently part of the AT&T Video Manager development team, he is involved in developing the technology for interactive video networking. Mr. Midani's current research interests are in connection control and management of BISDN networks and signaling for the Interactive Video Network. He received a B.S. from Wayne State University, Detroit, Michigan, and an M.S. from Arizona State University, Tempe, both in electrical engineering. Mr. Midani joined AT&T in 1986.*

Irene J. Morgan *is a distinguished member of technical staff in the Video Multimedia Department at AT&T Bell Laboratories in Lisle, Illinois. Currently part of the AT&T Video Manager systems engineering team, she is involved in defining the Video Manager product for interactive video networks. Ms. Morgan received a B.S. in computer engineering from the University of Michigan, Ann Arbor, and an M.S. in computer science from Columbia University, New York, New York. Ms. Morgan joined AT&T in 1980.*

Hung Tan Nguyen *is a distinguished member of technical staff in the AT&T Resource Link Organization in Lisle, Illinois. He is currently part of the AT&T Video Manager development team involved in developing the technology for interactive video networks. He is also doing research in concurrent programming languages for his doctoral degree in computer sciences at the Illinois Institute of Technology, Chicago. Mr. Nguyen, who received a B.S. in electrical engineering from the University of Illinois, Chicago, and an M.S. in computer sciences from the Illinois Institute of Technology, joined AT&T in 1984.*