

DAVIC and Interoperability Experiments

*Hari Kalva, Shih-Fu Chang, and Alexandros Eleftheriadis,
{hari, sfchang, eleft}@ee.columbia.edu*

Image and Advanced Television Laboratory,
Department of Electrical Engineering,
Columbia University, New York, NY 10027, USA

Abstract: To promote the development and deployment of broadband audio visual services, DAVIC has been working on specifying the standards and interfaces for end-to-end systems providing services to consumers over broadband networks. DAVIC specifies tools to provide these broadband services. Since the tools specified by DAVIC are standardized/developed by independent organizations, it is important to verify that the complete system works when these tools are used together. An interoperability event was organized by Columbia University for this very purpose. In this paper we first introduce DAVIC and DAVIC systems and then give details of the interoperability event: the motivation, goals, and results.

Key words: DAVIC, multimedia standardization, interoperability

1. Introduction to DAVIC

The Digital Audio Visual Council (DAVIC) is a consortium of companies, research institutes, and universities standardizing tools for end-to-end systems to provide broadband audio visual services [1, 2]. DAVIC was established in June 1994 and following an aggressive schedule, released the first specification, DAVIC 1.0 in December 1995. DAVIC 1.0 specifies the tools for the basic DAVIC system. DAVIC is continuing its work by specifying extensions to the specification. The first extension, DAVIC 1.1, was released in September 1996, and the second extension, DAVIC 1.2, will be released in December 1996. DAVIC for the most part picks existing standards for its specifications. However, if a standard does not exist for a particular technology, DAVIC creates a new standard. The following table shows some of the technologies selected by DAVIC for the 1.0 specification.

Table 1: Technologies Chosen in DAVIC 1.0

Tool	Technology
Video	MPEG-2, MP@ML (ISO/IEC 13818-2)
Audio	MPEG-1 (ISO/IEC 11172-3)
Core network	ATM
Access network	FTTH, FTTC, broadcast, ADSL, HFC...
User control	DSMCC-UU (ISO 13818-6)
Session control	DSMCC-UU (ISO 13818-6)
Connection control	Q.2931 (CCITT/ITU-T Q.2931)

DAVIC systems are expected to support a wide range of applications. DAVIC has identified 19 core applications and prioritized them based on interest. Some of the intended applications are,

movies on demand, tele-shopping, broadcast, games, karaoke on demand, tele-work, and news on demand.

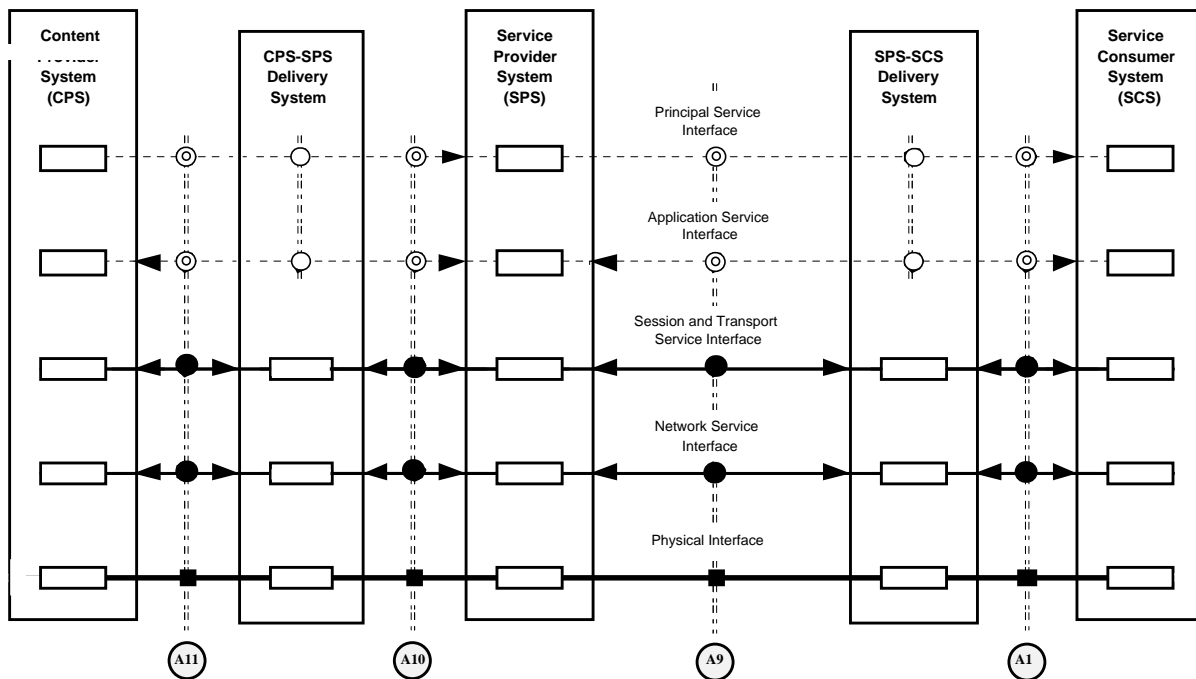


Figure 1. DAVIC Reference Model [1]

Figure 1 shows the reference model of a DAVIC system. The figure shows the reference points in the system and the interfaces between the sub-systems at these reference points. Reference points assume importance in a DAVIC system as a compliant implementation needs to be compliant only at these reference points. The information flow between the sub-systems is divided into five flows, S1 through S5. S1 information flow corresponds to the principal service information, e.g., MPEG-2 transport stream carrying audio and video. The S2 flow corresponds to the user control information, S3 corresponds to session control, S4 to connection management, and S5 corresponds to billing and other management functions. The S1 flow is unidirectional from the service provider to the service consumer (STU) in DAVIC 1.0 while the other information flows are bidirectional. Future versions of the specification are expected to include a bidirectional S1 flow supporting applications such as video conferencing. The five logical information flows may use a single physical channel or more than one physical channel. In the following sub-sections, a brief description of the components of a DAVIC system is given.

1.1. Service Consumer System (SCS)

SCS is another terminology for customer premise equipment (CPE) and usually consists of a NIU, an STU, a display terminal such as a TV, and other peripherals. The SCS is expected to include an in-house network and the ability to connect various peripherals in future extensions of the specification. SCS is connected to the Service Provider System (SPS) by means of a Delivery System (DS). Several interfaces were specified within a SCS, but only a few of them are defined

in DAVIC 1.0. The STU should be able to decode MPEG-2 transport streams with one or more programs. The transport stream packets are mapped to ATM cells as specified by the ATM Forum with the restriction that two transport packet be mapped to eight ATM cells. The upper limit on the MPEG-2 video bitrate is 15 Mbps (Main Profile at Main Level). For compressed audio MPEG-1 audio is specified and MHEG 5 is specified for encoding multimedia objects. DAVIC does not specify how a SCS (STU) should be implemented. Only reference points in a SCS are specified and a compliant system should meet these specifications. Figure 2 shows the modules commonly found in the STUs and the interfaces specified by DAVIC between the modules.

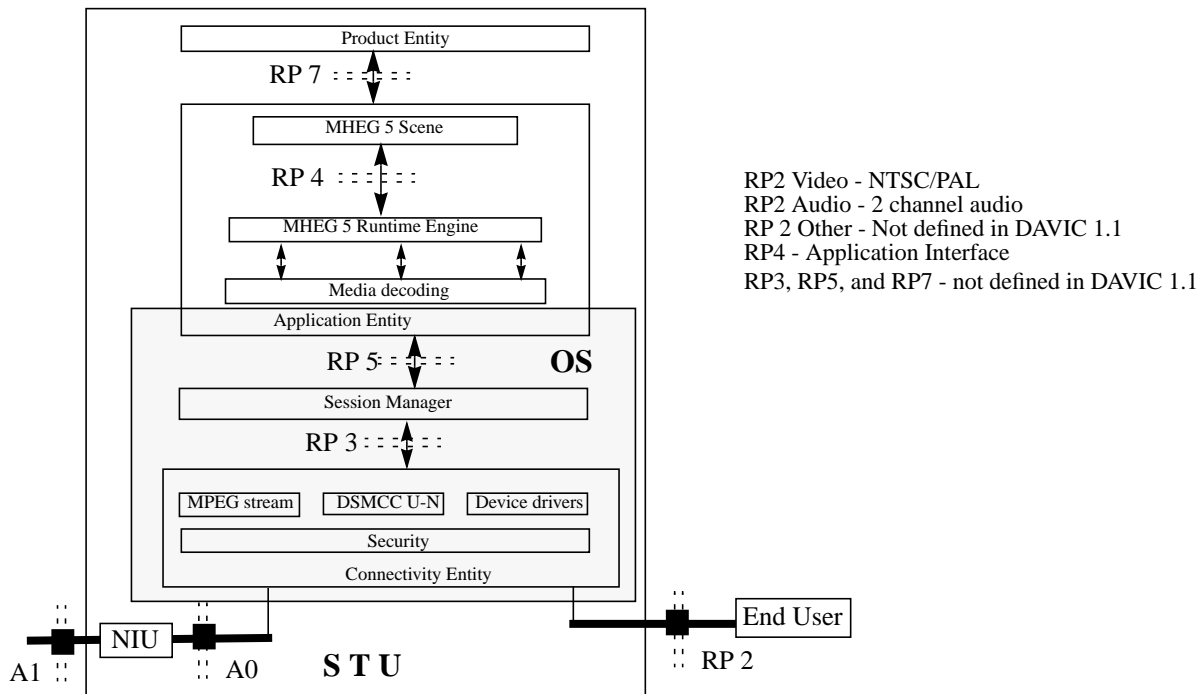


Figure 2. Common Modules in an STU

In response to the growing popularity of Internet-based services, the extension DAVIC 1.1 includes additional tools in an STU to provide Internet services to the consumer. A Java virtual machine is used to provide Internet services. An external data port to the STU to connect a PC is also specified. In DAVIC 1.1 this is an Ethernet port. The Reference Decoder Model (RDM) is also specified to enable application developers to develop applications that execute on STUs. Figure 3 shows the RDM specified in DAVIC 1.1. This model is an extension of the MPEG-2 Transport System Target Decoder model (T-STD). Only the DAVIC extensions to the T-STD are shown in the figure. The data enters the transport buffer (TB_n) from an MPEG-2 TS. The raw content used by the application is transferred to the content buffer (B_{contents}) at a rate R_x(n) specified by the application. The coded MHEG 5 objects are transferred to a code buffer (B_{code}). B_{code} also contains the Java VM bytecode. The run-time engine executes the code in B_{code} by loading it in to the memory (B_{execute}) and controls the times at which contents are decoded and removed from B_{contents}. Every DAVIC application includes a resource descriptor which provides information about the memory requirements, resolution and number of colors, and the transfer rate R_x(n), the rate of transfer from TB_n to B_{contents} and B_{code}. Every DAVIC 1.1 compliant STU shall have a minimum buffer size of 128 KB for B_{code}, 512 KB for B_{contents}, and 128 KB for

Bexecute.

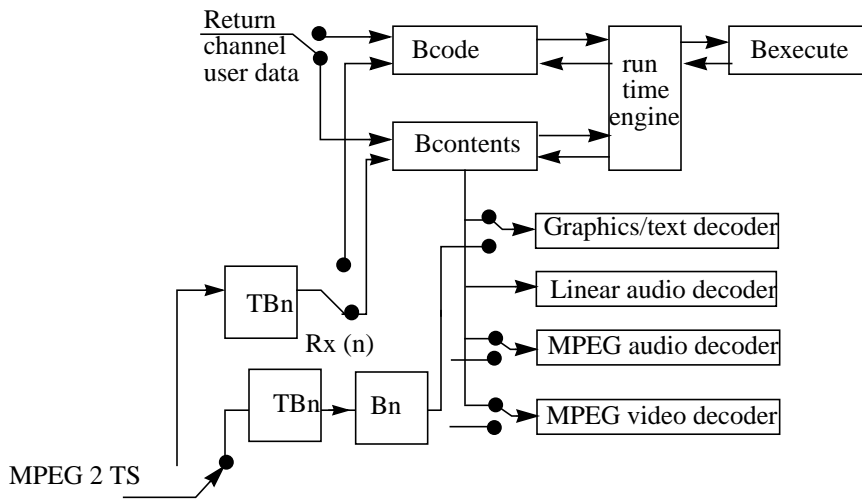


Figure 3. DAVIC Reference Decoder Model

1.2. Delivery System

Delivery System is a means to deliver information from a service provider to the service consumer. The means can be networked or non-networked (CD-ROM, tape etc.). In DAVIC 1.0 and 1.1 only networked delivery systems are considered. As shown in Figure 1, the service provider and service consumer are connected by a delivery system and so are content provider and service provider. However, in DAVIC 1.0/1.1, only the delivery system between A9 and A1 reference points is considered.

The delivery system is composed of a core network and an access network. The core network is ATM based and for the access network several access types are supported. There is a defined reference point, A4, between the core network and the access network. For Hertzian access networks (satellite/terrestrial broadcast), the return channel over PSTN/ISDN networks to support interactivity is specified in the 1.1 extension.

DAVIC has responded to the growing popularity of the Internet by specifying tools for accessing Internet-based services using the STU. The tools specified in the 1.1 extension to achieve this are, STU data port to connect a PC to the STU, cable modem, and Java virtual machine. DAVIC Internet access provides the end user with 'full' or 'client' Internet access using a PC connected to the STU. The different levels of Internet access are defined in [3]. Figure 4 shows the protocol stack for Internet access with upstream traffic for all cases and for end-to-end ATM downstream traffic. Figure 5 (a) shows the protocol stack used for down stream IP traffic when ATM terminates in the access network. In this case, IP packets are encapsulated in MPEG 2 transport stream packets. For IP over ATM, classical IP over ATM AAL 5 as specified in RFC 1483 and 1577 is used.

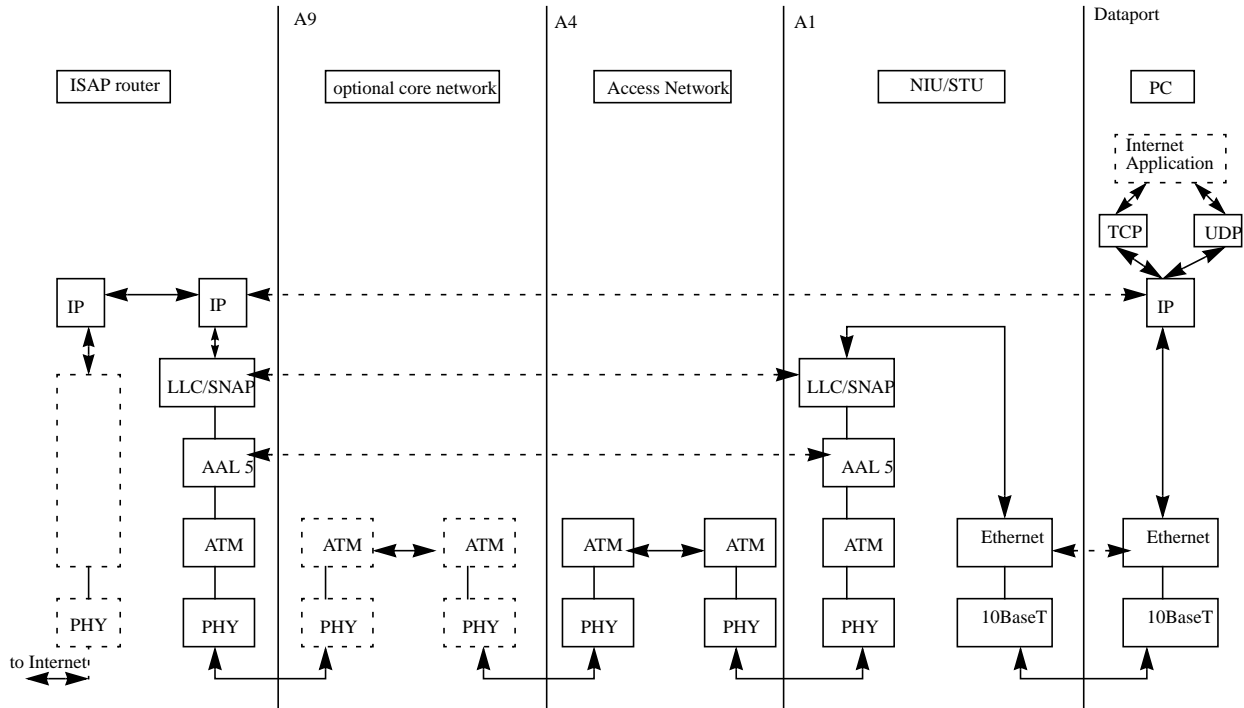


Figure 4. Internet Access: Upstream - all cases; Downstream - ATM end-to-end

For Internet access, Internet Access Service Provider (ISAP) is pre-agreed between the network provider and the end user. A PC is connected to the STU dataport as shown in Figure 5.b. In DAVIC 1.1 support of Ethernet on the PC-STU connection is normative [4]. Connection to the ISAP may be a pre-established through ATM PVCs or dynamically established using Q.2931 signaling. IP addresses may be pre-assigned or dynamically obtained using protocols such as DHCP. STU contains a data forwarder module which forwards the IP packets from the dataport (PC) to the network element. This module is responsible for LLC/SNAP encapsulation to and from the ISAP.

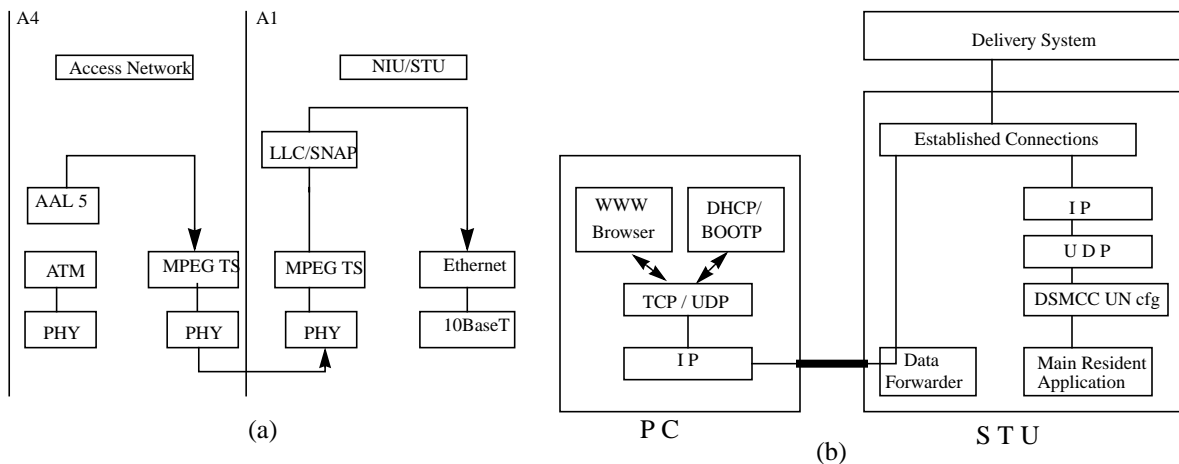


Figure 5. Internet Access

1.3. Service Provider System

SPS provides DAVIC services such as movies on demand. Unlike the SCS components in a DAVIC system, SPS does not have internal reference points. Only an external reference point A9, between SPS and the core network, is specified. The server reference model specified is a object oriented distributed server model with a collection of well defined services. Four core services are specified in the server model: service gateway, application service, stream service, and content service.

Service gateway is the STU's entry point in a server. Services available from a SPS are registered with the service gateway and de-registered when a service is discontinued. The service gateway allows an STU to browse and select service. It is also responsible for session management including resource provisioning and re-provisioning for a session. The application service provides a set of basic functions required by all the services such as directory functions. Additional services are built from application services by adding service-specific functions. The stream service is the primary one and is responsible for storing and delivering media streams. Content service is used by information providers to install new services and content on the server. Additional services such as session gateway, file, client profile, and download service are also specified.

2. An end-to-end Scenario in a DAVIC System

In this section, an end-to-end scenario including server selection, connection setup, session setup, application selection, content selection, and content playback is described. Before establishing a session, a client must first be able to locate a service provider. This is done using a level 1 gateway (L1GW). L1GW makes all the SPSs accessible to a end user. When the user selects a SPS, L1GW establishes a connection and drops out. When a connection is initially established, default resources are allocated for the session. During session establishment, the DSMCC-download protocol may be used to download the client application from the server. The STU browses the server and selects a service. Depending on the bandwidth requirements of the service, resources for the downstream channel are re-provisioned. The user launches the interactive application (e.g., a movie or a video game). When the user quits, session is tore down and resources are released.

3. DAVIC Interoperability

To achieve its goal of global deployment of DAVIC systems and services, DAVIC places special emphasis on the interoperability of the sub-systems. Interoperability promotes global competition and speeds up the deployment of services to consumers. Since the DAVIC specification is made up of a number of independently developed standards, it is important to ensure that these standards work when used together in a DAVIC system. The work related to verifying the specification and its interoperability is undertaken by the Interoperability sub-group of DAVIC. The charter of the Interoperability sub-group is to verify that the specification does not have any 'gaps' and to provide guidelines for conformance testing. To promote the Interoperability efforts, DAVIC is promoting a series of interoperability events, the first of which took place at Columbia University in New York in June 1996 [5].

The New York *Interoperability Event* was held in parallel with the 13th DAVIC meeting in New York, hosted by the Image Technology for New Media Center (ITNM) [6] at Columbia University. It was the first multi-national DAVIC interoperability event organized and had the goal of testing and verifying the interfaces and protocols specified in DAVIC 1.0. A total of eight organizations from around the world participated in the event, interconnecting components (video servers, set-top-boxes, and applications) of different DAVIC-compliant video on demand systems with each other, using Columbia University's ATM network. The participating organizations were: Columbia University(USA), CSELT (Italy), DeTeBerkom/GMD FOKUS (Germany), GCL (Japan), GTE Labs (USA), Hewlett Packard IDACOM (Canada), NIST (USA) and NTT (Japan).

3.1. Interoperability Configuration

The Interoperability demonstrations were shown for three days, June 18 and June 21 during the DAVIC meeting and on June 24 during the *Video on Demand Workshop*¹ [7] sponsored by the ITNM center. Over 350 people attended these demonstrations. The Interoperability testing took place in three phases during May-August 1996. Extensive testing was also done in addition to these demos.

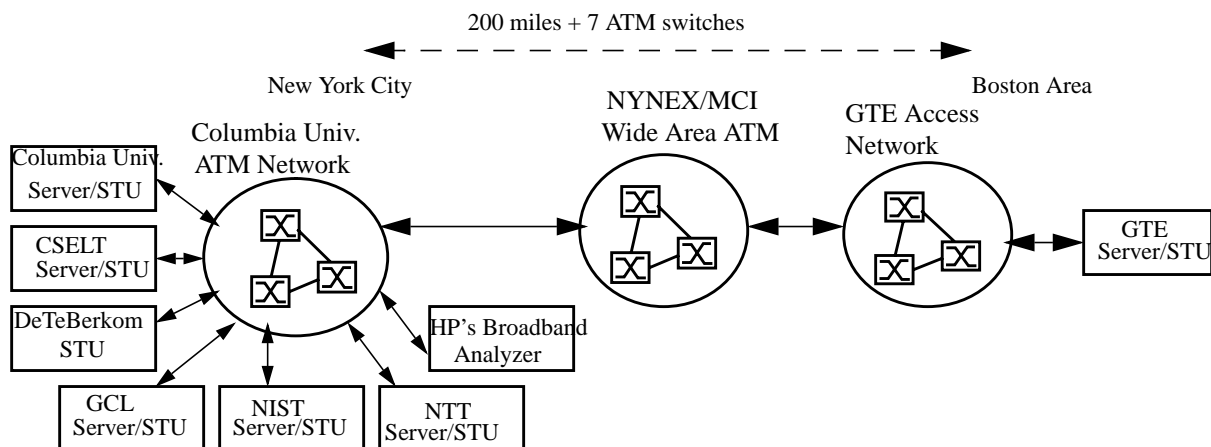


Figure 6. Interoperability Test Configuration

Figure 6 shows the configuration used in the Interoperability experiments. All the participants except GTE setup their VOD systems in the labs of the ITNM center. GTE was connected to the ITNM center via a wide area ATM network. As shown in the figure, each participating organization had a client and a server connected to Columbia's ATM network. DeTeBerkom participated with a client. HP participated with test equipment which was used for measuring MPEG-2 transport stream parameters.

The subnet of Columbia's ATM network used for the experiments was composed of three ATM switches: a FORE ASX100, a FORE ASX200, and an NEC Model-5 switch. The physical connections used multi-mode fiber and co-axial cable. The wide area connection included a DS3 link provided by NYNEX and MCI. It also included two ATM switches in the MCI segment and two ATM switches in the GTE local network.

1. For a copy of the proceedings, please send an e-mail to hari@itnm.columbia.edu

The main emphasis of the experiments was to test S1 and S2 information flows at the A1 and A9 reference points as specified by DAVIC (see Fig. 1). A1 is the reference point between a SCS and the network while A9 is the reference point between the SPS and the network. The information flows are explained in section 1. There are several reasons for restricting the scope of the Interoperability experiments to S1 and S2 information flows. At the time the experiments were conducted, there were no fully compliant DAVIC systems available. As a DAVIC system is very complex, it takes significant time and effort to build a complete system and there was not enough time before the event to build a complete system. The systems that were demonstrated were research prototypes. In the development of these prototypes, efforts were prioritized according to the importance of the information flows and S1, the MPEG data flow, and S2, the user control of the S1 flow, were developed first. This approach of prioritizing the information flows tested was also reflected in the decision of the Interoperability sub-group to develop conformance guidelines for S1 and S2 flows first.

The protocols specified by DAVIC for S1 and S2 information flows at A1 and A9 reference points are listed in columns 1 and 2 of tables 2. For the unidirectional S1 information flow from a server to a client, DAVIC specified MPEG 2 (1) video and MPEG 1 audio. The MPEG data is packetized as packetized elementary streams (PES) and the PES packets are multiplexed into a MPEG 2 transport stream (TS) as specified in the MPEG 2 systems standard [8]. The transport stream packets are mapped to ATM AAL 5 PDUs with 2 TS packets mapped to 8 ATM cells. The physical layer specified by DAVIC is SONET. For the S2 bidirectional user control, DSMCC UU primitives are used [9]. The primitives are encoded using OMG's Common Data Representation (CDR) and the RPC used to invoke the control commands is OMG's UNO. For the transport layer TCP is specified. The IP packets are mapped to ATM using classical IP over ATM as specified in RFC 1577.

All the systems were fully DAVIC compliant for S1 flow at A1 and A9 interfaces i.e. protocols as specified by DAVIC for S1 information flow were implemented by all the systems. At the physical layer, DAVIC specifies SONET, while some participants used TAXI. This does not create problems as the ATM switch converts between the formats. For the S2 flow, NTT and GCL are fully DAVIC compliant. All other participants are DAVIC-compliant except in the RPC used (XDR/ORBIX is used instead of CDR/OMG UNO). As expected this caused interoperability problems the details of which are given in Section 3.3.

In the following tables, the first column contains the DAVIC tools for that information flow. In the cells corresponding to each organization, a YES indicates DAVIC compliance and in case where non-compliant, the alternative tool used is indicated.

Table 2: Information Flows at A1 and A9

S 1 Flow	S 2 Flow
MPEG data	DSM-CC UU
PES	OMG-CDR
MPEG-2 TS	OMG-UNO
AAL 5	TCP
ATM	IP

Table 2: Information Flows at A1 and A9

S 1 Flow	S 2 Flow
SONET	AAL 5
-	ATM
-	SONET

3.2. S1 Interoperability

S1 interoperability is the ability of a system to decode and display the S1 information delivered by another system. S1 interoperability was tested among all the participants except DeTeBerkom. DeTeBerkom took part only in S2 interoperability testing with CSELT. Since connection and session signaling was not used, to connect company A's server to company B's client, a PVC was established between these components. Also, since there was limited S2 interoperability among the participants (see section 3.3 below), an MPEG-2 transport stream (TS) was delivered to the specific PVC without going through content selection by the client. Thus to test S1 interoperability between Server A and client B, the server A delivers the stream to PVC *ab* and client B decodes the stream arriving on PVC *ab*. The S1 Interoperability testing verified the interoperability of MPEG-2 TS generated by individual organizations. For the purpose of this experiments, Columbia generated several TSs each with different encoding parameters such as PCR (program clock reference) frequency and PSI (program specific information) table frequency. Even though all these streams are compliant to the MPEG-2 systems standard, some of them were not decoded well by few decoders. The reason for this is different interpretation and use of these values by the decoders. NTT also tested the bitmaps and graphics delivery as specified by DAVIC. CSELT's and GCL's clients were able to decode both NTSC and PAL streams. Since the CSELT server had MPEG-2 video sequences encoded from PAL source, there were some problems in interoperating with it. DAVIC specifies that when an STU receives a 525/29.97 (625/25) encoded signal it produce a 525/29.97 (625/25) output at the RP2 video interface (see Fig. 2) [10]. This limits the usage of encoded content to the regions in which the source format is used. Further more, a global content producer is forced to maintain a copy of encoded stream for each TV standard. To overcome this STUs should be able to produce the output suitable for the local display system (TV) irrespective of the source format.

3.3. S2 Interoperability

S2 information flow is a bidirectional flow used to transmit user control messages. If two systems have interoperable S2 flows, they can access each other's server and invoke DSMCC UU primitives to select and play video streams. S2 interoperability was tested between CSELT - DeTeBerkom, CSELT - Columbia, Columbia - NIST, and Columbia - GTE. This was possible as these companies have implemented the draft international standard (DIS) version of the DSMCC UU specification and use the same protocol stack for the S2 flow. For S2 interoperability testing between CSELT and DeTeBerkom, the test scenario consisted of DeTeBerkom's client retrieving MHEG5 objects from CSELT server and decoding the MHEG5 objects and controlling the streams using the DSM-CC primitives. In the S2 testing done between Columbia's client and CSELT server, file and stream objects (MPEG-2 TS) were accessed by Columbia's client from CSELT server. Columbia and NIST were able to demonstrate full S2 interoperability, including

server access, content selection and MPEG-2 stream playback.

3.4. Test Methodology and Results

Testing was mainly based on observation. A broadband MPEG-2 analyzer from HP was also used to measure PCR interval and PCR jitter. To test S1 information flow, an MPEG-2 TS was delivered to a client and the quality of audio and video was observed. In case of problems, causes were explored. A cross-connection from company A's server to company B's client is considered working (interoperable) if video and audio were played back without any disruptions. The results of S1 and S2 flow interoperability among the participants are summarized in Table 3.

In Table 3, for a cell corresponding to A's server and B's client, the following conventions are used:

W	works without problems
w	works when pumping TS generated by B but problems encountered when TS generated by A was used
blank cell	scenario was not tested
m	works with minor problems
D	does not work

There were two reasons for untested scenarios (blank cells in the table). In the case involving GTE the reason was that the wide area ATM connection connecting GTE's server and client to Columbia University was not available during the first two phases when most of the testing was done. In the case involving NIST, the reason was that NIST's client had a software decoder which can decode bitrates of less than 1.5 Mbps and only Columbia's server had such low bitrate TS.

Table 3: S2 Interoperability Testing Summary

Server	client		Columbia		CSELT		DeTeBerkom		GCL		GTE		NIST		NNT	
	S1	S2	S1	S2	S1	S2	S1	S2	S1	Ss2	S1	S2	S1	S2		
Information Flow	W	W	m	m			w		W	W	W	W	w			
Columbia	w	m	W	W		W	W						w			
CSELT	w		W				W	W					w			
GCL	w		W				W	W					w			
GTE	w								W	W						
NIST	w	w	m				w				W	W	w			
NTT	w		m				D						w	w		

Since the systems tested were research prototypes, we cannot attribute the interoperability problems totally to 'gaps' in the DAVIC specification. Some of the problems faced could be due to implementation problems. The S1 interoperability problems between CSELT and other participants were due to the restriction in the DAVIC specification as explained in section 3.2. In some cases the decoded video was jerky. This depends on the jitter tolerance of the client system. Since all the servers used were software based running on general purpose operating systems, the output of the server cannot be guaranteed to be constant bit rate. This results in packet jitter at the output of the server which could contribute to jerky video in some decoders depending on the decoder's

jitter tolerance.

As shown in Table 3, the S2 flow was implemented by all organizations but interoperability could not be achieved among the majority of the participants. This is mainly because of the incompatible RPC mechanism used and the level of compliance to the DSMCC UU specification. It should be noted that the applications tested were simple, use a small percentage of the DSMCC UU primitives, and the servers are accessed by one user at a time for all the participating servers except Columbia's server which allows multiple user access. It is not clear how the RPC specified by DAVIC (OMG UNO) scales when the services are commercially deployed and accessed by thousands of users at the same time.

4. Conclusion

DAVIC is playing an important role in the global deployment of broadband audio-visual services. Currently DAVIC has successfully released its first specification 1.0 and also its first extension 1.1. Work is being continued to extend the specification to provide advanced services to consumers. Verifying the specification through interoperability experiments increases the credibility of the specification and encourages the industry to manufacture DAVIC compliant systems. The interoperability event provided an opportunity for all the participants to test and verify their implementation of the DAVIC specification. S1 flow interoperability was thoroughly tested while S2 flow interoperability was not completely tested. From these experiments, it can be concluded that the incompatibilities were mainly in content generation. There were some unsettled mismatches between some MPEG 2 transport streams and some decoders. This could be due to different interpretation and handling of timing information by the decoders. Continued interoperability testing is necessary to identify the problems in the DAVIC specifications. The second in a series of DAVIC interoperability events was held during the Tokyo Electronic Show in October 1996. This event also includes S3 flow interoperability in addition to S1 and S2 flows.

Acknowledgments

Acknowledgments are due to the other members of Columbia's VOD team [11] for their contributions to the testbed. Acknowledgments are also due to the organizations and their representatives (primary contact listed below) that took part in the event.

Organization	Contact Person	E-Mail
CSELT, Italy	Stefano Dal Lago	stefano.dallago@cse.lt.stet.it
DeTe Telekom, Germany	Hans Werner Bitzer	bitzer@deteteberkom.de
GCL, Japan	Shigeru Komatsu	komat@gctech.co.jp
GTE Labs, USA	Len Ulbricht	lwu1@gte.com
HP IDACOM, Canada	Trevor Dyck	trevord@idacom.hp.com
NIST, USA	David Hsu	su@isdn.ncsl.nist.gov
NTT, Japan	Hisashi Kasahara	kasahara@nttvdt.hil.ntt.jp

Columbia University acknowledges C-Cube, GTE Labs, HP-IDACOM, IBM, NEC, NYNEX, and SGI for their support for the testbed infrastructure.

References

- [1] The Digital Audio-Visual Council (DAVIC), "Part 2 - Systems Reference Models and Scenarios," DAVIC 1.0 specification, Geneva, Switzerland, 1995.
- [2] A. Eleftheriadis, "Architecting Video-on-Demand Systems: DAVIC 1.0 and Beyond", Proceedings, International Symposium on Multimedia Communications and Video Coding, Brooklyn, NY, October 1995
- [3] D. Crocker, "To be on the Internet," RFC 1775, IETF, March 1995.
- [4] The Digital Audio-Visual Council (DAVIC), "Internet Access Tools," Baseline document No. 28, DAVIC 1.1 specification, 1996.
- [5] Columbia University, "DAVIC New York Interoperability Experiments: Report and Results," DAVIC Contribution No. DAVIC/TC/SYS/96/09/008, Geneva, Switzerland, 1996.
- [6] DAVIC VoD Interoperability Test Event, June 1996, Columbia University, New York (<http://www.itnm.columbia.edu/vod/interop96>)
- [7] VOD96, Workshop on Video on Demand: Technology, Interoperability, and Trials, URL: <http://www.itnm.columbia.edu/workshopo/vod96>.
- [8] Generic Coding of Moving Pictures and Associated Audio (MPEG 2 Systems), ISO/IEC 13818-1, International Standard, 1995.
- [9] Generic Coding of Moving Pictures and Associated Audio (DSMCC), ISO/IEC 13818-6, International Standard, 1996.
- [10] The Digital Audio-Visual Council (DAVIC), "Part 5 - Service Consumer System and High Level API," DAVIC 1.0 specification, 1995.
- [11] S.-F. Chang, A. Eleftheriadis, D. Anastassiou, S. Jacobs, H. Kalva, and J. Zamora, "Columbia's VoD and Multimedia Research Testbed With Heterogeneous Network Support," appears in this special issue, Intern. J. on Multimedia Tools and Applications, Kluwer Academic Publishers, May, 1997.