Editor: Peiya Liu
Siemens Corporate Research

MPEG-7

The Generic Multimedia Content Description Standard, Part 1

José M. Martínez
Polytechnic University of Madrid, Spain

Rob Koenen
InterTrust Technologies

Fernando Pereira
Technical University of Lisbon, Portugal

An immeasurable amount of multimedia information is available today—in digital archives, on the Web, in broadcast data streams, and in personal and professional databases—and this amount continues to grow. Yet, the value of that information depends on how easily we can manage, find, retrieve, access, and filter it.

The transition between two millennia abounds with new ways to produce, offer, filter, search, and manage digitized multimedia information. Broadband is being offered with increasing audio and video quality, using ever-improving access speeds on both fixed and mobile networks. As a result, users are confronted with numerous content sources. Wading through these sources, and finding what you need and what you like in the vast content sea, is becoming a daunting task.

On the other end of the usage spectrum, the task at hand is becoming equally challenging. For content providers who have many (digital and analog) sources, the effort to find the right piece of content increases with the amount available. The question of identifying and managing content isn't just restricted to database retrieval applications such as digital libraries but extends to areas like broadcast channel selection, multimedia editing, surveillance, and home-entertainment devices.

MPEG-7—developed by the Moving Picture Experts Group (MPEG)—addresses this content management challenge. (This same International Organization for Standardization [ISO] committee also developed the successful standards known as MPEG-1 [1992], MPEG-2 [1995], and MPEG-4 [version 1 in 1998 and version 2 in 1999].) The recently completed ISO/IEC International Standard 15938, formally called the Multimedia Content Description Interface (but better known as MPEG-7), provides a rich set of tools for completely describing multimedia content. The standard wasn't just designed from a content management viewpoint (classical archival information). It includes an innovative description of the media's content, which we can extract via content analysis and processing. MPEG-7 also isn't aimed at any one application; rather, the elements that MPEG-7 standardizes support as broad a range of applications as possible. This is one of the key differences between MPEG-7 and other metadata standards; it aims to be generic, not targeted to a specific application or application domain.

This article provides a comprehensive overview of MPEG-7's motivation, objectives, scope, and components. In the July–September 2002 issue of *IEEE MultiMedia*, we'll provide more details about the MPEG-7 description tools.

Objectives and scope

The different parts of the MPEG-7 standard collectively and individually offer a comprehensive set of multimedia description tools to create so-called *descriptions*, which can be used by applications that enable quality access to content. This implies good storage solutions; high-performance content identification; and fast, ergonomic, accurate, and personalized filtering, searching, and retrieval. MPEG-7 gives a generic framework that can support various applications. Ideally, it facilitates exchange and reuse of multimedia content across different application domains.

Although people traditionally use multimedia information (for instance, by watching a TV show or listening to music), in an increasing number of cases, computational systems create, exchange, retrieve, and reuse multimedia information. This is the case for scenarios that employ image understanding (such as surveillance, intelligent vision, and smart cameras) and media conversion (such as speech to text, picture to speech, and speech to picture). Other scenarios are information retrieval and filtering multiple streams of

1070-986X/02/$17.00 © 2002 IEEE

multimedia content using content descriptions (such as automatically selecting and recording a TV program with a hard-disk-based personal video recorder). Yet another example is an image sensor that triggers an alarm when a certain visual or audio event happens. Automatic transcoding may be performed from a string of characters to audible information, or a search may be performed in a stream of audio or video data.

Multimedia sources will play an increasingly pervasive role in our lives, and there's a growing need to have these sources processed further. This makes it necessary to develop standardized forms of multimedia information representation that go beyond those aimed at presentation, such as analog waveforms (PAL and NTSC video standards), digital samples (A-law or μ-law Pulse Code Modulation [PCM] audio), frame-based compressed (MPEG-1 and MPEG-2), or even object-based compressed representations (MPEG-4). These representation schemes have evolved to provide better ways for transmitting, storing, and retrieving multimedia information, ultimately for human consumption. Today, we need new forms of representation, allowing some degree of interpretation of the information's meaning, that a device or a computer code can access.

In the examples we just gave, an image sensor can produce visual data in the form of objects with associated physical measures and time information. A processing device can locally or remotely process these to verify if certain programmed conditions are met. A video recording device could receive descriptions of the multimedia information associated with a program that would enable it to record, for example, only news with the exclusion of sports. We could describe products from a company so that a machine could respond to unstructured queries from customers.

Early in the MPEG-7 requirements gathering process,[1] the MPEG's Requirements Group collected and studied several applications to extract requirements for the standard. MPEG-7 addresses applications that we can store (online or offline) or stream (for example, broadcast and push models on the Internet) and can operate in both real-time and nonreal-time environments. (A real-time environment in this context means that the description is generated while the content is being captured.) Examples of applications the MPEG's Requirements Group considered[2] include digital multimedia libraries, broadcast media selection, multimedia editing, home-entertainment devices, searching multimedia content, managing (large) content archives, opening up archives to the public, and semiautomatic multimedia presentation and editing.

A large number of descriptive features of multimedia content must be considered to cover all these applications, but for each application, application domain, and usage context, only a subset of descriptive features will be used. This is an important point because no single, correct description of a given body of content exists. Rather, many different possible descriptions exist, and all are equally valid. We can describe the same content using different types of features, depending on usage. Also, we can describe the features in different configurations. (We explain how this works later on.) MPEG-7 distinguishes itself from other relevant metadata standards in its support for a range of abstraction levels, from low-level signal characteristics to high-level semantic information. For example, for visual material, a lower abstraction level would be a description of a video object's shape, size, texture, color, movement (trajectory), and position. An audio description would include key, mood, tempo, tempo changes, and position in sound space. The highest level would give semantic information—This is a scene with a barking brown dog on the left and a blue ball that falls down on the right, with the sound of passing cars in the background. Intermediate abstraction levels can also exist, such as genre classification and content ratings.

The abstraction level relates to the way we extract the features. We can automatically extract most low-level features, whereas high-level features usually need human supervision and annotation. However, the science of image and sound understanding keeps advancing, and MPEG-7 can fully support this progress because only the description format is fixed and not the extraction methodologies.

In addition to describing features inherently

Further Reading

Find more information about MPEG-7 at the MPEG homepage (http://www.tilab.com/mpeg). This Web page contains links to a wealth of information about MPEG, including much about MPEG-7, many publicly available MPEG documents, several lists of frequently asked questions, and links to other MPEG-7 Web pages. The MPEG-7 Alliance Web site (http://www.mpeg-industry.com) contains information and links more related to applications and industry.

Neither automatic nor semiautomatic feature extraction algorithms are inside the scope of the standard because their standardization isn't required to allow interoperability.

in the content, MPEG-7 must also support the description of other types of information about the multimedia data such as form (for example, the coding scheme used or the overall data size), place and time of recording, classification, and links to other relevant material. In the MPEG-7 usage chain, first the descriptions are generated using the MPEG-7 format, then stored or streamed, and finally consumed in a way depending on the application in question.

Producing the description starts with feature extraction (via analysis of the content) or annotation. Automatically extracting features for fully exploiting the possibilities of MPEG-7 descriptions will be extremely useful. However, automatic extraction isn't always possible. The higher the abstraction level, the more difficult automatic extraction is, and interactive tools (for human supervision and annotation) will be of good use. However useful they may be, neither automatic nor semiautomatic feature extraction algorithms are inside the scope of the standard. The main reason is that their standardization isn't required to allow interoperability. Also, leaving this open means there's room for competition among different providers and thus for innovation. Moreover, not standardizing the analysis phase means the use of the standard can grow with expected technological innovation.

The MPEG-7 standard also doesn't specify the consumption end of the chain—search engines, filter agents, or any other description consuming entity. This isn't necessary, and again, competition and innovation will produce the best results. Moreover, it's impossible to predict the types of applications that will use the MPEG-7 Description format.

This means that only the Description format—the minimum needed to ensure interoperability among applications—is standardized, which is the standard's goal.

Structure

Before explaining the MPEG-7 standard's structure, we'll first mention the major MPEG-7 elements, which include Descriptors, Description Schemes, a Description Definition Language, and Systems Tools. (We used the definitions from the MPEG-7 Requirements Document[1] to define these elements.)

Data is multimedia "information that will be described using MPEG-7, regardless of storage, coding, display, transmission, medium, or technology." Furthermore, a Feature is "a distinctive characteristic of the data [that] signifies something to somebody." Based on these definitions, these are the four major MPEG-7 building blocks:

- *Descriptor*: "A representation of a Feature. A Descriptor defines the syntax and the semantics of the Feature representation."
- *Description Scheme*: "The structure and semantics of the relationships between its components, which may be both Descriptors and Description Schemes."
- *Description Definition Language* (DDL): "A language that allows the creation of new Description Schemes and, possibly, Descriptors. It also allows the extension and modification of existing Description Schemes."
- *Systems Tools*: Tools to support multiplexing of descriptions, synchronization of descriptions with content, delivery mechanisms, and coded representations (both textual and binary formats) for efficient storage and transmission and the management and protection of intellectual property in MPEG-7 Descriptions.

With these tools, we can build an MPEG-7 Description and deploy it. According to the requirements document,[1] "a Description consists of a Description Scheme (structure) and the set of Descriptor Values (instantiations) that describe the Data." A Descriptor Value is "an instantiation of a Descriptor for a given data set (or subset thereof)."

The MPEG-7 standard consists of several parts. This lets us use the various clusters of technology alone, according to MPEG's toolbox approach to standardization. It also keeps the editing of the standard manageable. Here are the eight MPEG-7 standard parts:

- *MPEG-7 Systems*[3] specifies the system tools to prepare MPEG-7 Descriptions for efficient transport and storage (through binarization) to allow synchronization between content and descriptions.
- *MPEG-7 Description Definition Language*[4] specifies the DDL.
- *MPEG-7 Visual*[5] specifies the description tools dealing with visual-only descriptions.
- *MPEG-7 Audio*[6] specifies the description tools dealing with audio-only descriptions. This includes both Descriptors and Description Schemes.
- *MPEG-7 Multimedia Description Schemes*[7] specifies the Descriptors and Description Schemes for generic features and multimedia descriptions. This includes the Descriptors and Description Schemes that aren't specific to audio or video.
- *MPEG-7 Reference Software*[8] gives a software implementation of the MPEG-7 standard parts we list here. The MPEG-7 Reference Software, which is based on the MPEG Experimentation Model, is normative with respect to the decoding behavior. This means that all conformant implementations should produce the same decoding results as the reference software. We can use this reference software free of copyright for building compliant MPEG-7 implementations. (Note that free of copyright isn't the same as free of patent rights.)
- *MPEG-7 Conformance*[9] specifies the guidelines and procedures for testing whether descriptions and consumption engines are compliant with the MPEG-7 standard.
- *MPEG-7 Extraction and Use of Descriptions* provides information on the extraction and use of some of the description tools. Notably, it gives insight into the reference software and alternative approaches.

MPEG-7 Descriptions

With the description tools in MPEG-7, we can create different types of Descriptions. Classical archival-oriented Descriptions include

- information regarding the content's creation and production processes (such as director, title, actors, and location),
- information related to using the content (such as broadcast schedules and copyright pointers), and
- information on storing and representing the content (such as storage format and encoding format).

A second example is innovative perceptual Descriptions of the information in the content:

- information regarding the content's spatial, temporal, or spatio–temporal structure (for example, scene cuts, segmentation in regions, and region motion tracking);
- information about low-level features in the content (for example, colors, textures, sound timbres, and melody descriptions); and
- semantic information related to the reality captured by the content (for example, objects, events, and interactions between objects).

The third type is additional information for organizing, managing, and accessing the content:

- information about how objects are related and gathered in collections,
- information to support efficient browsing content (such as summaries, variations, and transcoding information), and
- information about the interaction of the user with the content (such as user preferences and usage history).

These three types of Descriptions are interrelated, and we can combine them in a single Description and even in a single Description Scheme. (Remember that a Description is a set of instantiated Description Schemes.) MPEG-7 is unique in the breadth of the Descriptors the standard specifies, and it distinguishes itself with its

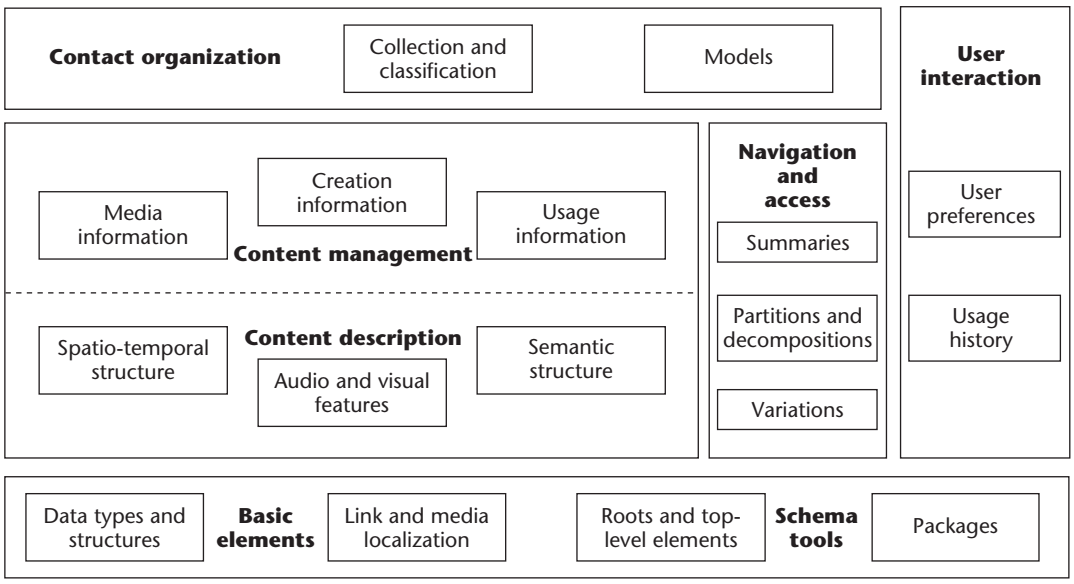


Figure 1. MPEG-7 description tools.

structuring capabilities. We can combine traditional metadata, low-level descriptors, and high-level semantic information in a single Description in a structured way, laying explicit links between these types of data features.

We can physically locate MPEG-7 Descriptions with the associated multimedia material, in the same data stream, or on the same storage system, or locate them somewhere else. When the content and its descriptions aren't colocated, mechanisms that link multimedia material and their MPEG-7 Descriptions are necessary.

The content data and the query data don't have to be of the same type. For example, we can query visual material using visual content, music, speech, and so forth. It's the responsibility of the search engine and filter agent to match the query data to the MPEG-7 Descriptions.

MPEG-7 tools

Here we give a brief description of the tools in the different parts of the MPEG-7 standard. Part 2 of this article (in the July–September issue) will provide a more complete description of the MPEG-7 description tools.

MPEG-7 description tools

The Descriptors and Description Schemes together form the set of MPEG-7's predefined description tools. We can group them in different classes according to their functionality (see Figure 1).

Basic Elements. The Basic Elements are the generic entities that are used as building blocks by various description tools. They include basic data types (such as numbers, matrices, vectors, and country); links and locators (such as time, media locators, and referencing tools); and other basic description tools for places, people, textual annotations, controlled vocabularies, and so on.

Schema Tools. The Schema Tools are the tools for wrapping description tools for use by applications. This group also includes the package tools for organizing related description tools into groups with personalized labels for easing the use by specific applications.

Content Description Tools. Content Description Tools represent perceptible information, including structural aspects (structure description tools), audio and visual features, and conceptual aspects (semantic description tools).

The Structure description tools let us describe content in terms of spatio–temporal segments organized in a hierarchical structure—for example, letting us define a table of contents or an index. We can attach audio, visual, annotation, and content management description tools to the segments to describe them in detail.

MPEG-7 Visual description tools include the visual basic structures (such as description tools for grid layout, time series, and spatial coordinates) and visual description tools that let us describe color, texture, shape, motion, localization, and faces.

MPEG-7 Audio description tools comprise the

audio description framework (including the scale tree for creating a scalable series of audio samples, low-level audio descriptors, and the silence descriptor) and high-level audio description tools that let us describe musical instrument timbre, sound recognition, spoken content, and melody.

The Semantic description tools let us describe the content with real-world semantics and conceptual notions: objects, events, abstract concepts, and relationships. We can cross-link the semantic and structure description tools with a set of links.

Content Management Tools. The Content Management Tools let us specify information about media features, creation, and usage of multimedia content.

The Media description tools let us describe the storage media, coding format, quality, and transcoding hints for adapting content to different networks and terminals.

The Creation description tools let us describe the creation process (for example, title, agents, materials, places, and dates), classification (for example, genre, subject, parental rating, and languages), and related materials.

The Usage description tools let us describe the conditions for use (for example, rights and availability) and the history of use (for example, financial results and audience).

Content Organization Tools. Content Organization Tools let us create and model collections of multimedia content and descriptions. We can describe each collection as a whole by their attribute values characterized by models and statistics.

Navigation and Access Tools. Navigation and Access Tools let us specify summaries, partitions and decompositions, and variations of multimedia content for facilitating browsing and retrieval.

Summary description tools provide both hierarchical and sequential navigation modes to provide efficient preview access to the multimedia material.

Partitions and decompositions description tools allow multiresolution and progressive access in time, space, and frequency.

Variations description tools let us describe preexisting views of multimedia content: summaries, different media modalities, (for example, image and text), scaled versions, and so on.

Adopting the XML Schema as the basis for the MPEG-7 DDL and the resulting XML-compliant instances eases interoperability by using a common, generic, and powerful representation format.

User Interaction Tools. User Interaction Tools let us describe user preferences (for personalized filtering) and usage history pertaining to the consumption of the multimedia content.

Using MPEG-7 description tools. The MPEG-7 description tools are a library of standardized Descriptors and Description Schemes. This library is presented on the basis of the functionality they provide, but in practice, we can combine them into meaningful sets of description units making use of the Schema tools. Each application builder might want to select a subset of Descriptors and Description Schemes. Currently, the discussion about how to deal with this selection in a normative way is still open within MPEG. The experience gained in profiling previous MPEG standards was a key asset in addressing this topic, and the lesson learned is that it's better to wait with creating normative subsets until there's real implementation experience in the industry. Some applications might need to incorporate specific description tools that aren't (yet) in the standard. We can do this using the DDL.

MPEG-7 DDL

The DDL lets us describe the MPEG-7 description tools (Descriptors and Description Schemes) and extend them with application specific description tools. Although in MPEG-7's early development, the DDL was to be a markup language fully defined within MPEG, fairly early on the DDL Ad Hoc Group decided to change course and develop it on top of the Extensible Markup

Figure 2. MPEG-7 Description (in textual format) of this article.

```
<Mpeg7 xmlns="http://www.mpeg7.org/2001/MPEG-7_Schema" xml:lang="en"
type="complete">
<ContentManagement xsi:type="CreationInformationType">
   <Creation>
           <Title xml:lang="en">
           MPEG-7: the generic Multimedia Content Description Standard
           </Title>
           <Abstract>
                 <FreeTextAnnotation>
                              An overview of the MPEG-7 standard.
                 </FreeTextAnnotation>
                 <StructuredAnnotation>
                       <What><Name>MPEG-7 Overview</Name></What>
      </StructuredAnnotation>
           </Abstract>
           <Creator>
                 <Role href="urn:mpeg:mpeg7:cs:RoleCS:AUTHOR">
                 <Agent xsi:type="PersonType">
                       <Name>
                              <GivenName>Rob</GivenName>
                              <FamilyName>Koenen</FamilyName>
                       </Name>
                 </Agent>
           </Creator>
           <Creator>
                 <Role href="urn:mpeg:mpeg7:cs:RoleCS:AUTHOR">
                 <Agent xsi:type="PersonType">
                       <Name>
                              <GivenName>Fernando</GivenName>
                              <FamilyName>Pereira</FamilyName>
                       </Name>
                 </Agent>
           </Creator>
           <CreationCoordinates>
                 <CreationDate>
                       <TimePoint>2001-07-13</TimePoint>
                       <Duration>P7D</Duration>
                 </CreationDate>
           </CreationCoordinates>
   </Creation>
   <Classification>
           <Genre href="urn:mpeg:ContentCS:1">
                 <Name xml:lang="en">Information</Name>
           </Genre>
           <Language type="main">en</Language>
           <Release date="2001-09"/>
   </Classification>
   <RelatedMaterial>
           <MediaLocator>
                 <MediaUri>http://www.tilab.com/mpeg/</MediaUri>
           </MediaLocator>
   </RelatedMaterial>
</ContentManagement>
</Mpeg7>
```

```
<Mpeg7 xmlns="http://www.mpeg7.org/2001/MPEG-7_Schema" xml:lang="en"
type="complete">
<ContentDescription xsi:type="ContentEntityType">
<MultimediaContent xsi:type="ImageType">
<Image>
       <MediaLocator>
              <MediaUri>
              http://www.tilab.org/mpeg/mpeg_logo-anim_1.gif
              </MediaUri>
       </MediaLocator>
       <CreationInformation">
              <Creation>
                    <Title xml:lang="en">The animated MPEG Logo</Title>
                    <Creator>
                          <Role href="urn:mpeg:mpeg7:cs:RoleCS:AUTHOR">
                                <Name xml:lang="en">Author</Name>
                          </Role>
                          <Agent xsi:type="OrganizationType">
                                <Name>MPEG</Name>
                          </Agent>
                    </Creator>
              </Creation>
              <RelatedMaterial>
                    <MediaLocator>
                          <MediaUri>http://www.tilab.com/mpeg/</MediaUri>
                    </MediaLocator>
              </RelatedMaterial>
       </CreationInformation>
</Image>
</MultimediaContent>
</ContentDescription>
</Mpeg7>
```

(a)

(b)

Figure 3. (a) MPEG-7 Description (in textual format) of the MPEG logo. (b) The MPEG logo.

Language (XML) Schema.[10–12] The DDL adds extensions to the XML Schema to address specific MPEG-7 requirements such as support for vectors, matrices, and typed references.[1]

Adopting the XML Schema as the basis for the MPEG-7 DDL and the resulting XML-compliant instances (which constitute Descriptions in MPEG-7 textual format; see Figures 2 and 3) eases interoperability by using a common, generic, and powerful representation format. It also facilitates the adoption of MPEG-7 tools for extending existing XML applications with multimedia content description functionalities. (NewsML, XML for News, is one example of this. Find more information at http://www.iptc.org.)

MPEG-7 Systems Tools

MPEG-7 Systems partially include the neces-

MPEG-7 addresses many different applications in many different environments, which means it needs to provide a flexible and extensible framework for describing multimedia data.

sary tools to prepare MPEG-7 Descriptions for efficient transport and storage, to allow synchronization between descriptions and between content and descriptions, and to support dynamic update and incremental build-up of Descriptions in a receiving terminal.

The generic concept of delivery covers the transport, storage, and retrieval functionalities. This delivery aspect is outside the scope of the standard, as is the case in MPEG-4. Like MPEG-2 and MPEG-4, MPEG-7 uses the concept of *elementary streams*. A delivery layer (which is not defined in MPEG-7) provides MPEG-7 elementary streams to the decompression layer. MPEG-7 elementary streams consist of consecutive individually accessible portions of data named *access units*—another concept inherited from MPEG-2 and MPEG-4. Elementary streams contain Schema and Description information. The decompression layer is in charge of parsing the access units, reconstructing the content description, and passing it to the application.

We can represent Descriptions in text format, binary format, or using a mixture of the two, depending on the usage scenario. The text format is defined with the DDL and the instances are XML files. The binary format (BiM, or Binary Format for MPEG-7 data) provides descriptions in an alternative, dual representation. The binary descriptions have a more efficient compression than the textual descriptions and provide flexibility for streaming and efficient random search and access to description elements without the need of parsing the complete bitstream. The binary format supports the binary representation of the standardized description tools but currently not the extension to application-specific description tools as the DDL does.

Conclusions and future work

MPEG-7 is the standard for describing multimedia content, providing the richest multimedia description tools for content management, organization, navigation, and automated processing. The MPEG-7 standard defines a large library of core description tools and a DDL for those who want to extend the standard. A set of Systems Tools provides the means for deploying the descriptions in specific storage and transport environments. Because MPEG-7 is a generic standard, not all the description tools are necessary for all applications.

As we've already mentioned, MPEG-7 addresses many different applications in many different environments, which means it needs to provide a flexible and extensible framework for describing multimedia data. With this in mind, MPEG-7 was designed to take into account relevant work by other leading standards such as the Society for Motion Piction and Television Engineers Metadata Dictionary (SMPTE), Dublin Core, European Broadcasting Union (EBU) P/Meta, and TV Anytime (which are all focused on more or less specific applications or application domains).

MPEG-7 became ISO/IEC 15398 standard in fall 2001, but the MPEG committee is already working on future extensions. Possible additions under discussion are description tools for describing synthetic content and interactivity with content and linguistic tools. **MM**

Acknowledgments

This article is based on the MPEG public document *Overview of the MPEG-7 Standard*.[13] We thank all the contributors to this living document.

References

1. *ISO/MPEG N4320, MPEG-7 Requirements Document*, v 15, F. Pereira, ed., MPEG Requirements Group, Sydney, July 2001.
2. *ISO/MPEG N4676, MPEG-7 Applications*, v 11.0, N. Day, ed., MPEG Requirements Group, Jeju, Mar. 2002.
3. *ISO/MPEG N4285, Text of ISO/IEC Final Draft International Standard 15938-1 Information Technology -*

Multimedia Content Description Interface - Part 1 Systems, MPEG Systems Group, Sydney, July 2001.
4. *ISO/MPEG N4288, Text of ISO/IEC Final Draft International Standard 15938-2 Information Technology - Multimedia Content Description Interface - Part 2 Description Definition Language*, MPEG Systems Group, Sydney, July 2001.
5. *ISO/MPEG N4358, Text of ISO/IEC Final Draft International Standard 15938-3 Information Technology - Multimedia Content Description Interface - Part 3 Visual*, MPEG Video Group, Sydney, July 2001.
6. *ISO/MPEG N4224, Text of ISO/IEC Final Draft International Standard 15938-4 Information Technology - Multimedia Content Description Interface - Part 4 Audio*, MPEG Audio Group, Sydney, July 2001.
7. *ISO/MPEG N4242, Text of ISO/IEC Final Draft International Standard 15938-5 Information Technology - Multimedia Content Description Interface - Part 5 Multimedia Description Schemes*, MPEG Multimedia Description Schemes Group, Sydney, July 2001.
8. *ISO/MPEG N4206, Text of ISO/IEC Final Draft International Standard 15938-6 Information Technology - Multimedia Content Description Interface - Part 6 Reference Software*, MPEG Implementation Studies Group, Sydney, July 2001.
9. *ISO/MPEG N4633, FCD 15938-7 Information Technology-Multimedia Content Description Interface-Part 7 Conformance*, MPEG Committee, Jeju, Mar. 2002.
10. *XML Schema Part 0: Premier*, World Wide Web Consortium Recommendation, May 2001, http://www.w3.org/TR/xmlschema-0.
11. *XML Schema Part 1: Structures*, World Wide Web Consortium Recommendation, May 2001, http://www.w3.org/TR/xmlschema-1.
12. *XML Schema Part 2: Datatypes*, World Wide Web Consortium Recommendation, May 2001, http://www.w3.org/TR/xmlschema-2.
13. *ISO/MPEG N4674, Overview of the MPEG-7 Standard*, v 6.0, J.M. Martínez, ed., MPEG Requirements Group, Jeju, Mar. 2002.

Readers may contact José M. Martínez at the Image Processing Group, Polytechnic University of Madrid, Depto. Señales, Sistemas y Radiocomunicaciones, E.T.S.Ing. Telecomunicación, Universidad Politécnica de Madrid, Ciudad, E-28040 Madrid, Spain, email jms@gti.ssr.upm.es.

Contact Standards editor Peiya Lui, Siemens Corporate Research, 755 College Rd. East, Princeton, NJ 08540, email Peiya.Liu@scr.siemens.com

Media Reviews continued from p. 77

However, even though the book achieves its stated goal of providing a reference for a wide variety of technologies, I felt disappointed by the apparent lack of cohesion. There was no organized way to locate certain specific information. Reading sequentially from beginning to end is a laborious process, given the range of topics that this book covers. The book jumps quickly from simple fundamental topics—something at an introductory undergraduate level—to important topical high-level issues.

Given the intended audience and Watkinson's quick shifts from basic to advanced technological ideology, there's danger of losing novice and expert audiences. It would have been helpful for him to provide some markers to identify the intended level of the reader throughout the book. Also, the author's presentation of personal philosophies (throughout the text) as fundamental dogmas diminishes the value of this book as an authoritative reference. For example, in the introduction, the author states that the use of region codes in DVDs is "exactly the wrong solution" (p. 36). Similar observations were made about commercial television, such as "the business model is a crude one which evolved when technology was equally crude" (p. 33). I feel that such personal viewpoints don't belong in a reference book.

However, using the book selectively, I'd recommend ignoring the author's personal viewpoints and skipping over sections based on the reader's prior understanding of the relevant technologies. Also, I think that specific targeted books on the various technologies will complement and enhance the information covered in this book. **MM**

Readers may contact Surendar Chandra or Media Reviews editor Amit Sheth at the Computer Science Department, University of Georgia, 415 Boyd Graduate Studies Research Center, Athens, GA 30602-7404, email surendar@.cs.uga.edu or amit@cs.uga.edu.