
**Information technology — Coding of
audio-visual objects —**

**Part 10:
Advanced Video Coding**

**AMENDMENT 2: MVC extensions for
inclusion of depth maps**

Technologies de l'information — Codage des objets audiovisuels —

Partie 10: Codage visuel avancé

*AMENDMENT 2: Extensions du codage vidéo multivues pour
l'inclusion de cartes de profondeur*





COPYRIGHT PROTECTED DOCUMENT

© ISO 2013

All rights reserved. Unless otherwise specified, no part of this publication may be reproduced or utilized otherwise in any form or by any means, electronic or mechanical, including photocopying, or posting on the internet or an intranet, without prior written permission. Permission can be requested from either ISO at the address below or ISO's member body in the country of the requester.

ISO copyright office
Case postale 56 • CH-1211 Geneva 20
Tel. + 41 22 749 01 11
Fax + 41 22 749 09 47
E-mail copyright@iso.org
Web www.iso.org

Published in Switzerland

Foreword

ISO (the International Organization for Standardization) and IEC (the International Electrotechnical Commission) form the specialized system for worldwide standardization. National bodies that are members of ISO or IEC participate in the development of International Standards through technical committees established by the respective organization to deal with particular fields of technical activity. ISO and IEC technical committees collaborate in fields of mutual interest. Other international organizations, governmental and non-governmental, in liaison with ISO and IEC, also take part in the work. In the field of information technology, ISO and IEC have established a joint technical committee, ISO/IEC JTC 1.

International Standards are drafted in accordance with the rules given in the ISO/IEC Directives, Part 2.

The main task of the joint technical committee is to prepare International Standards. Draft International Standards adopted by the joint technical committee are circulated to national bodies for voting. Publication as an International Standard requires approval by at least 75 % of the national bodies casting a vote.

Attention is drawn to the possibility that some of the elements of this document may be the subject of patent rights. ISO and IEC shall not be held responsible for identifying any or all such patent rights.

Amendment 2 to ISO/IEC 14496-10:2012 was prepared by Joint Technical Committee ISO/IEC JTC 1, *Information technology*, Subcommittee SC 29, *Coding of audio, picture, multimedia and hypermedia information*.

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC 14496-10:2012/Amd 2:2013

Information technology — Coding of audio-visual objects —

Part 10:

Advanced Video Coding

AMENDMENT 2: MVC extensions for inclusion of depth maps

In 0.6, add the following paragraph after the paragraph that starts with "Multiview video coding":

An extension of multiview video coding that additionally supports the inclusion of depth maps is specified in Annex I, allowing the construction of bitstreams that represent multiple views with corresponding depth views. In a similar manner as with the multiview video coding specified in Annex H, bitstreams encoded as specified in Annex I may also contain sub-bitstreams that conform to this Specification.

In 0.7, add the following paragraph after the paragraph that starts with "Annex H specifies":

Annex I specifies MVC extensions for inclusion of depth maps, referred to as multiview video coding with depth (MVCD). The reader is referred to Annex I for the entire decoding process for MVCD, which is specified there with references being made to clauses 2-9 and Annexes A-E and Annex H. Subclause I.10 specifies one profile for MVCD (Multiview and Depth).

In Clause 2, add the following additional normative reference:

- ISO 12232:2006, *Photography – Digital still cameras – Determination of exposure index, ISO speed ratings, standard output sensitivity, and recommended exposure index.*

In Clause 4, add the following additional abbreviation:

MVCD Multiview Video Coding with Depth

In 7.3.1, replace the syntax table with:

nal_unit(NumBytesInNALunit) {	C	Descriptor
forbidden_zero_bit	All	f(1)
nal_ref_idc	All	u(2)
nal_unit_type	All	u(5)
NumBytesInRBSP = 0		
nalUnitHeaderBytes = 1		
if(nal_unit_type == 14 nal_unit_type == 20 nal_unit_type == 21) {		
svc_extension_flag	All	u(1)
if(svc_extension_flag)		
nal_unit_header_svc_extension() /* specified in Annex G */	All	
else		
nal_unit_header_mvc_extension() /* specified in Annex H */	All	
nalUnitHeaderBytes += 3		
}		
for(i = nalUnitHeaderBytes; i < NumBytesInNALunit; i++) {		
if(i + 2 < NumBytesInNALunit && next_bits(24) == 0x000003) {		
rbsp_byte [NumBytesInRBSP++]	All	b(8)
rbsp_byte [NumBytesInRBSP++]	All	b(8)
i += 2		
emulation_prevention_three_byte /* equal to 0x03 */	All	f(8)
} else		
rbsp_byte [NumBytesInRBSP++]	All	b(8)
}		
}		

In 7.3.2.1.1, replace the syntax table with:

seq_parameter_set_data() {	C	Descriptor
profile_idc	0	u(8)
constraint_set0_flag	0	u(1)
constraint_set1_flag	0	u(1)
constraint_set2_flag	0	u(1)
constraint_set3_flag	0	u(1)
constraint_set4_flag	0	u(1)
constraint_set5_flag	0	u(1)
reserved_zero_2bits /* equal to 0 */	0	u(2)
level_idc	0	u(8)
seq_parameter_set_id	0	ue(v)
if(profile_idc == 100 profile_idc == 110 profile_idc == 122 profile_idc == 244 profile_idc == 44 profile_idc == 83 profile_idc == 86 profile_idc == 118 profile_idc == 128 profile_idc == 138) {		
chroma_format_idc	0	ue(v)
if(chroma_format_idc == 3)		
separate_colour_plane_flag	0	u(1)
bit_depth_luma_minus8	0	ue(v)
bit_depth_chroma_minus8	0	ue(v)
qpprime_y_zero_transform_bypass_flag	0	u(1)
seq_scaling_matrix_present_flag	0	u(1)
if(seq_scaling_matrix_present_flag)		
for(i = 0; i < ((chroma_format_idc != 3) ? 8 : 12); i++) {		
seq_scaling_list_present_flag[i]	0	u(1)
if(seq_scaling_list_present_flag[i])		
if(i < 6)		
scaling_list(ScalingList4x4[i], 16, UseDefaultScalingMatrix4x4Flag[i])	0	
else		
scaling_list(ScalingList8x8[i - 6], 64, UseDefaultScalingMatrix8x8Flag[i - 6])	0	
}		
}		
log2_max_frame_num_minus4	0	ue(v)
pic_order_cnt_type	0	ue(v)
if(pic_order_cnt_type == 0)		
log2_max_pic_order_cnt_lsb_minus4	0	ue(v)
else if(pic_order_cnt_type == 1) {		
delta_pic_order_always_zero_flag	0	u(1)
offset_for_non_ref_pic	0	se(v)
offset_for_top_to_bottom_field	0	se(v)
num_ref_frames_in_pic_order_cnt_cycle	0	ue(v)
for(i = 0; i < num_ref_frames_in_pic_order_cnt_cycle; i++)		
offset_for_ref_frame[i]	0	se(v)
}		
max_num_ref_frames	0	ue(v)
gaps_in_frame_num_value_allowed_flag	0	u(1)

pic_width_in_mbs_minus1	0	ue(v)
pic_height_in_map_units_minus1	0	ue(v)
frame_mbs_only_flag	0	u(1)
if(!frame_mbs_only_flag)		
mb_adaptive_frame_field_flag	0	u(1)
direct_8x8_inference_flag	0	u(1)
frame_cropping_flag	0	u(1)
if(frame_cropping_flag) {		
frame_crop_left_offset	0	ue(v)
frame_crop_right_offset	0	ue(v)
frame_crop_top_offset	0	ue(v)
frame_crop_bottom_offset	0	ue(v)
}		
vui_parameters_present_flag	0	u(1)
if(vui_parameters_present_flag)		
vui_parameters()	0	
}		

In 7.3.2.1.3, replace the syntax table with:

subset_seq_parameter_set_rbsp() {	C	Descriptor
seq_parameter_set_data()	0	
if(profile_idc == 83 profile_idc == 86) {		
seq_parameter_set_svc_extension() /* specified in Annex G */	0	
svc_vui_parameters_present_flag	0	u(1)
if(svc_vui_parameters_present_flag == 1)		
svc_vui_parameters_extension() /* specified in Annex G */	0	
} else if(profile_idc == 118 profile_idc == 128) {		
bit_equal_to_one /* equal to 1 */	0	f(1)
seq_parameter_set_mvc_extension() /* specified in Annex H */	0	
mvc_vui_parameters_present_flag	0	u(1)
if(mvc_vui_parameters_present_flag == 1)		
mvc_vui_parameters_extension() /* specified in Annex H */	0	
} else if(profile_idc == 138) {		
bit_equal_to_one /* equal to 1 */	0	f(1)
seq_parameter_set_mvcd_extension() /* specified in Annex I */	0	
}		
additional_extension2_flag	0	u(1)
if(additional_extension2_flag == 1)		
while(more_rbsp_data())		
additional_extension2_data_flag	0	u(1)
rbsp_trailing_bits()	0	
}		

In 7.3.3, replace the syntax table with:

slice_header() {	C	Descriptor
first_mb_in_slice	2	ue(v)
slice_type	2	ue(v)
pic_parameter_set_id	2	ue(v)
if(separate_colour_plane_flag == 1)		
colour_plane_id	2	u(2)
frame_num	2	u(v)
if(!frame_mbs_only_flag) {		
field_pic_flag	2	u(1)
if(field_pic_flag)		
bottom_field_flag	2	u(1)
}		
if(IdrPicFlag)		
idr_pic_id	2	ue(v)
if(pic_order_cnt_type == 0) {		
pic_order_cnt_lsb	2	u(v)
if(bottom_field_pic_order_in_frame_present_flag && !field_pic_flag)		
delta_pic_order_cnt_bottom	2	se(v)
}		
if(pic_order_cnt_type == 1 && !delta_pic_order_always_zero_flag) {		
delta_pic_order_cnt[0]	2	se(v)
if(bottom_field_pic_order_in_frame_present_flag && !field_pic_flag)		
delta_pic_order_cnt[1]	2	se(v)
}		
if(redundant_pic_cnt_present_flag)		
redundant_pic_cnt	2	ue(v)
if(slice_type == B)		
direct_spatial_mv_pred_flag	2	u(1)
if(slice_type == P slice_type == SP slice_type == B) {		
num_ref_idx_active_override_flag	2	u(1)
if(num_ref_idx_active_override_flag) {		
num_ref_idx_l0_active_minus1	2	ue(v)
if(slice_type == B)		
num_ref_idx_l1_active_minus1	2	ue(v)
}		
if(nal_unit_type == 20 nal_unit_type == 21)		
ref_pic_list_mvc_modification() /* specified in Annex H */	2	
else		
ref_pic_list_modification()	2	
if((weighted_pred_flag && (slice_type == P slice_type == SP)) (weighted_bipred_idc == 1 && slice_type == B))		
pred_weight_table()	2	
if(nal_ref_idc != 0)		
dec_ref_pic_marking()	2	
if(entropy_coding_mode_flag && slice_type != I && slice_type != SI)		
cabac_init_idc	2	ue(v)
slice_qp_delta	2	se(v)

if(slice_type == SP slice_type == SI) {		
if(slice_type == SP)		
sp_for_switch_flag	2	u(1)
slice_qs_delta	2	se(v)
}		
if(deblocking_filter_control_present_flag) {		
disable_deblocking_filter_idc	2	ue(v)
if(disable_deblocking_filter_idc != 1) {		
slice_alpha_c0_offset_div2	2	se(v)
slice_beta_offset_div2	2	se(v)
}		
}		
if(num_slice_groups_minus1 > 0 && slice_group_map_type >= 3 && slice_group_map_type <= 5)		
slice_group_change_cycle	2	u(v)
}		

Replace Table 7-1 with:

nal_unit_type	Content of NAL unit and RBSP syntax structure	C	Annex A NAL unit type class	Annex G and Annex H NAL unit type class	Annex I NAL unit type class
0	Unspecified		non-VCL	non-VCL	non-VCL
1	Coded slice of a non-IDR picture slice_layer_without_partitioning_rbsp()	2, 3, 4	VCL	VCL	VCL
2	Coded slice data partition A slice_data_partition_a_layer_rbsp()	2	VCL	not applicable	not applicable
3	Coded slice data partition B slice_data_partition_b_layer_rbsp()	3	VCL	not applicable	not applicable
4	Coded slice data partition C slice_data_partition_c_layer_rbsp()	4	VCL	not applicable	not applicable
5	Coded slice of an IDR picture slice_layer_without_partitioning_rbsp()	2, 3	VCL	VCL	VCL
6	Supplemental enhancement information (SEI) sei_rbsp()	5	non-VCL	non-VCL	non-VCL
7	Sequence parameter set seq_parameter_set_rbsp()	0	non-VCL	non-VCL	non-VCL
8	Picture parameter set pic_parameter_set_rbsp()	1	non-VCL	non-VCL	non-VCL
9	Access unit delimiter access_unit_delimiter_rbsp()	6	non-VCL	non-VCL	non-VCL
10	End of sequence end_of_seq_rbsp()	7	non-VCL	non-VCL	non-VCL
11	End of stream end_of_stream_rbsp()	8	non-VCL	non-VCL	non-VCL
12	Filler data filler_data_rbsp()	9	non-VCL	non-VCL	non-VCL
13	Sequence parameter set extension seq_parameter_set_extension_rbsp()	10	non-VCL	non-VCL	non-VCL
14	Prefix NAL unit prefix_nal_unit_rbsp()	2	non-VCL	suffix dependent	suffix dependent
15	Subset sequence parameter set subset_seq_parameter_set_rbsp()	0	non-VCL	non-VCL	non-VCL
16..18	Reserved		non-VCL	non-VCL	non-VCL
19	Coded slice of an auxiliary coded picture without partitioning slice_layer_without_partitioning_rbsp()	2, 3, 4	non-VCL	non-VCL	non-VCL
20	Coded slice extension slice_layer_extension_rbsp()	2, 3, 4	non-VCL	VCL	VCL
21	Coded slice extension for depth view components /*specified in Annex I */ slice_layer_extension_rbsp() /*	2, 3, 4	non-VCL	non-VCL	VCL

	specified in Annex I */				
22..23	Reserved		non-VCL	non-VCL	VCL
24..31	Unspecified		non-VCL	non-VCL	non-VCL

In 7.4.1, make the following changes:

Replace the following:

svc_extension_flag indicates whether a `nal_unit_header_svc_extension()` or `nal_unit_header_mvc_extension()` will follow next in the syntax structure.

with:

svc_extension_flag indicates whether a `nal_unit_header_svc_extension()` or `nal_unit_header_mvc_extension()` will follow next in the syntax structure. When `nal_unit_type` is equal to 21, `svc_extension_flag` shall be equal to 0 and the semantics of `svc_extension_flag` equal to 1 are reserved for future specification by ITU-T | ISO/IEC.

Add the following paragraph after the semantics of `svc_extension_flag` just before the semantics of `rbbsp_bytef[i]`.

The value of `svc_extension_flag` shall be equal to 0 for coded video sequences conforming to one or more profiles specified in Annex I. Decoders conforming to one or more profiles specified in Annex I shall ignore (remove from the bitstream and discard) NAL units for which `nal_unit_type` is equal to 14, 20, or 21 and `svc_extension_flag` is equal to 1.

In 7.4.2.1.1, replace the following:

chroma_format_idc specifies the chroma sampling relative to the luma sampling as specified in clause 6.2. The value of `chroma_format_idc` shall be in the range of 0 to 3, inclusive. When `chroma_format_idc` is not present, it shall be inferred to be equal to 1 (4:2:0 chroma format).

with

chroma_format_idc specifies the chroma sampling relative to the luma sampling as specified in clause 6.2. The value of `chroma_format_idc` shall be in the range of 0 to 3, inclusive. When `chroma_format_idc` is not present and `profile_idc` is not equal to 138, `chroma_format_idc` shall be inferred to be equal to 1 (4:2:0 chroma format). When `chroma_format_idc` is not present and `profile_idc` is equal to 138, `chroma_format_idc` shall be inferred to be equal to 0 (4:0:0 chroma format), otherwise, it shall be inferred to be equal to 1 (4:2:0 chroma format).

In 7.4.2.1.3, replace the following:

additional_extension2_data_flag may have any value. It shall not affect the conformance to profiles specified in Annex A, G, or H.

with

additional_extension2_data_flag may have any value. It shall not affect the conformance to profiles specified in Annex A, G, H, or I.

Replace Annex C with:

Annex C Hypothetical reference decoder

(This annex forms an integral part of this Recommendation | International Standard)

This annex specifies the hypothetical reference decoder (HRD) and its use to check bitstream and decoder conformance.

Two types of bitstreams are subject to HRD conformance checking for this Recommendation | International Standard. The first such type of bitstream, called Type I bitstream, is a NAL unit stream containing only the VCL NAL units and filler data NAL units for all access units in the bitstream. The second type of bitstream, called a Type II bitstream, contains, in addition to the VCL NAL units and filler data NAL units for all access units in the bitstream, at least one of the following:

- additional non-VCL NAL units other than filler data NAL units,
- all `leading_zero_8bits`, `zero_byte`, `start_code_prefix_one_3bytes`, and `trailing_zero_8bits` syntax elements that form a byte stream from the NAL unit stream (as specified in Annex B).

Figure C-1 shows the types of bitstream conformance points checked by the HRD.




Figure C-1 – Structure of byte streams and NAL unit streams for HRD conformance checks

The syntax elements of non-VCL NAL units (or their default values for some of the syntax elements), required for the HRD, are specified in the semantics subclauses of clause 7, Annexes D and E, and subclauses G.7, G.13, G.14, H.7, H.13, H.14, I.7, I.13, and I.14.

Two types of HRD parameter sets (NAL HRD parameters and VCL HRD parameters) are used. The HRD parameter sets are signalled as follows:

- When the coded video sequence conforms to one or more of the profiles specified in Annex A and the decoding process specified in clauses 2-9 is applied, the HRD parameter sets are signalled through video usability information as specified in subclauses E.1 and E.2, which is part of the sequence parameter set syntax structure.
- When the coded video sequence conforms to one or more of the profiles specified in Annex G and the decoding process specified in Annex G is applied, the HRD parameter sets are signalled through the SVC video usability information extension as specified in subclauses G.14.1 and G.14.2, which is part of the subset sequence parameter set syntax structure.

NOTE 1 – For coded video sequences that conform to both, one or more of the profiles specified in Annex A and one or more of the profiles specified in Annex G, the signalling of the applicable HRD parameter sets is depending on whether the decoding process specified in clauses 2-9 or the decoding process specified in Annex G is applied.

- When the coded video sequence conforms to one or more of the profiles specified in Annex H and the decoding process specified in Annex H is applied, the HRD parameter sets are signalled through the MVC video usability information extension as specified in subclauses H.14.1 and H.14.2, which is part of the subset sequence parameter set syntax structure.

NOTE 2 – For coded video sequences that conform to both, one or more of the profiles specified in Annex A and one or more of the profiles specified in Annex H, the signalling of the applicable HRD parameter sets is depending on whether the decoding process specified in clauses 2-9 or the decoding process specified in Annex H is applied.

- When the coded video sequence conforms to one or more of the profiles specified in Annex I and the decoding process specified in Annex I is applied, the HRD parameter sets are signalled through the MVC video usability information extension as specified in subclause I.14, which is part of the subset sequence parameter set syntax structure.

NOTE 3 – For coded video sequences that conform to one or more of the profiles specified in Annex A, one or more of the profiles specified in Annex H and one or more of the profiles specified in Annex I, the signalling of the applicable HRD parameter sets is depending on whether the decoding process specified in clauses 2-9, the decoding process specified in Annex H or the decoding process specified in Annex I is applied.

All sequence parameter sets and picture parameter sets referred to in the VCL NAL units, and corresponding buffering period and picture timing SEI messages shall be conveyed to the HRD, in a timely manner, either in the bitstream (by non-VCL NAL units), or by other means not specified in this Recommendation | International Standard.

In Annexes C, D, and E and subclauses G.12, G.13, G.14, H.12, H.13, H.14, I.12, I.13, and I.14, the specification for "presence" of non-VCL NAL units is also satisfied when those NAL units (or just some of them) are conveyed to decoders (or to the HRD) by other means not specified by this Recommendation | International Standard. For the purpose of counting bits, only the appropriate bits that are actually present in the bitstream are counted.

NOTE 4 – As an example, synchronization of a non-VCL NAL unit, conveyed by means other than presence in the bitstream, with the NAL units that are present in the bitstream, can be achieved by indicating two points in the bitstream, between which the non-VCL NAL unit would have been present in the bitstream, had the encoder decided to convey it in the bitstream.

When the content of a non-VCL NAL unit is conveyed for the application by some means other than presence within the bitstream, the representation of the content of the non-VCL NAL unit is not required to use the same syntax specified in this annex.

NOTE 5 – When HRD information is contained within the bitstream, it is possible to verify the conformance of a bitstream to the requirements of this subclause based solely on information contained in the bitstream. When the HRD information is not present in the bitstream, as is the case for all "stand-alone" Type I bitstreams, conformance can only be verified when the HRD data is supplied by some other means not specified in this Recommendation | International Standard.

The HRD contains a coded picture buffer (CPB), an instantaneous decoding process, a decoded picture buffer (DPB), and output cropping as shown in Figure C-2.




Figure C-2 – HRD buffer model

The CPB size (number of bits) is $CpbSize[SchedSelIdx]$. The DPB size (number of frame buffers) is $Max(1, max_dec_frame_buffering)$. When the coded video sequence conforms to one or more of the profiles specified in Annex H and the decoding process specified in Annex H is applied, the DPB size is specified in units of view components. When the coded video sequence conforms to one or more of the profiles specified in Annex I and the decoding process specified in Annex I is applied, the DPB is operated separately for texture view components and depth view components and the terms texture DPB and depth DPB are used, respectively. The texture DPB size is specified in units of texture view components and the depth DPB size is specified in units of depth view components.

The HRD operates as follows. Data associated with access units that flow into the CPB according to a specified arrival schedule are delivered by the HSS. The data associated with each access unit are removed and decoded instantaneously by the instantaneous decoding process at CPB removal times. Each decoded picture is placed in the DPB at its CPB removal time unless it is output at its CPB removal time and is a non-reference picture. When a picture is placed in the DPB it is removed from the DPB at the later of the DPB output time or the time that it is marked as "unused for reference".

For each picture in the bitstream, the variable `OutputFlag` for the decoded picture and, when applicable, the reference base picture is set as follows:

- If the coded video sequence containing the picture conforms to one or more of the profiles specified in Annex A and the decoding process specified in clauses 2-9 is applied, `OutputFlag` is set equal to 1.
- Otherwise, if the coded video sequence containing the picture conforms to one or more of the profiles specified in Annex G and the decoding process specified in Annex G is applied, the following applies:
 - For a reference base picture, `OutputFlag` is set equal to 0.

- For a decoded picture, OutputFlag is set equal to the value of the output_flag syntax element of the target layer representation.
- Otherwise, if the coded video sequence containing the picture conforms to one or more of the profiles specified in Annex H and the decoding process specified in Annex H is applied, the following applies:
- For the decoded view components of the target output views, OutputFlag is set equal to 1.
- For the decoded view components of other views, OutputFlag is set equal to 0.
- Otherwise (the coded video sequence containing the picture conforms to one or more of the profiles specified in Annex I and the decoding process specified in Annex I is applied), the following applies:
 - For the decoded texture view components and corresponding depth view components with same VOIdx of the target output views, OutputFlag is set equal to 1.
 - For the decoded texture view components and corresponding depth view components with same VOIdx of other views, OutputFlag is set equal to 0.

The operation of the CPB is specified in subclause C.1. The instantaneous decoder operation is specified in clauses 2-9 (for coded video sequences conforming to one or more of the profiles specified in Annex A) and in Annex G (for coded video sequences conforming to one or more of the profiles specified in Annex G) and in Annex H (for coded video sequences conforming to one or more of the profiles specified in Annex H) and in Annex I (for coded video sequences conforming to one or more of the profiles specified in Annex I). The operation of the DPB is specified in subclause C.2. The output cropping is specified in subclause C.2.2.

NOTE 6 – Coded video sequences that conform to both, one or more of the profiles specified in Annex A and one or more of the profiles specified in Annex G, can be decoded either by the decoding process specified in clauses 2-9 or by the decoding process specified in Annex G. The decoding result and the HRD operation may be dependent on which of the decoding processes is applied.

NOTE 7 – Coded video sequences that conform both to one or more of the profiles specified in Annex A and one or more of the profiles specified in Annex H can be decoded either by the decoding process specified in clauses 2-9 or by the decoding process specified in Annex H. The decoding result and the HRD operation may be dependent on which of the decoding processes is applied.

NOTE 8 – Coded video sequences that conform to one or more of the profiles specified in Annex A, one or more of the profiles specified in Annex H and one or more of the profiles specified in Annex I, can be decoded either by the decoding process specified in clauses 2-9, by the decoding process specified in Annex H or by the decoding process specified in Annex I. The decoding result and the HRD operation may be dependent on which of the decoding processes is applied.

HSS and HRD information concerning the number of enumerated delivery schedules and their associated bit rates and buffer sizes is specified in subclauses E.1.1, E.1.2, E.2.1, E.2.2, G.14.1, G.14.2, H.14.1, H.14.2 and I.14. The HRD is initialised as specified by the buffering period SEI message as specified in subclauses D.1.1 and D.2.1. The removal timing of access units from the CPB and output timing from the DPB are specified in the picture timing SEI message as specified in subclauses D.1.2 and D.2.2. All timing information relating to a specific access unit shall arrive prior to the CPB removal time of the access unit.

When the coded video sequence conforms to one or more of the profiles specified in Annex G and the decoding process specified in Annex G is applied, the following is specified:

- (a) When an access unit contains one or more buffering period SEI messages that are included in scalable nesting SEI messages and are associated with values of DQId in the range of $((DQIdMax \gg 4) \ll 4)$ to $((DQIdMax \gg 4) \ll 4) + 15$, inclusive, the last of these buffering period SEI messages in decoding order is the buffering period SEI message that initialises the HRD. Let hrdDQId be the largest value of $16 * sei_dependency_id[i] + sei_quality_id[i]$ that is associated with the scalable nesting SEI message

containing the buffering period SEI message that initialises the HRD, let hrdDId and hrdQId be equal to $\text{hrdDQId} \gg 4$ and $\text{hrdDQId} \& 15$, respectively, and let hrdTId be the value of sei_temporal_id that is associated with the scalable nesting SEI message containing the buffering period SEI message that initialises the HRD.

- (b) The picture timing SEI messages that specify the removal timing of access units from the CPB and output timing from the DPB are the picture timing SEI messages that are included in scalable nesting SEI messages associated with values of $\text{sei_dependency_id}[i]$, $\text{sei_quality_id}[i]$, and sei_temporal_id equal to hrdDId , hrdQId , and hrdTId , respectively.
- (c) The HRD parameters that are used for conformance checking are the HRD parameters included in the SVC video usability information extension of the active SVC sequence parameter set that are associated with values of $\text{vui_ext_dependency_id}[i]$, $\text{vui_ext_quality_id}[i]$, and $\text{vui_ext_temporal_id}[i]$ equal to hrdDId , hrdQId , and hrdTId , respectively. For the specification in this annex, num_units_in_tick , time_scale , $\text{fixed_frame_rate_flag}$, $\text{nal_hrd_parameters_present_flag}$, $\text{vcl_hrd_parameters_present_flag}$, $\text{low_delay_hrd_flag}$, and $\text{pic_struct_present_flag}$ are substituted with the values of $\text{vui_ext_num_units_in_tick}[i]$, $\text{vui_ext_time_scale}[i]$, $\text{vui_ext_fixed_frame_rate_flag}[i]$, $\text{vui_ext_nal_hrd_parameters_present_flag}[i]$, $\text{vui_ext_vcl_hrd_parameters_present_flag}[i]$, $\text{vui_ext_low_delay_hrd_flag}[i]$, and $\text{vui_ext_pic_struct_present_flag}[i]$, respectively, with i being the value for which $\text{vui_ext_dependency_id}[i]$, $\text{vui_ext_quality_id}[i]$, and $\text{vui_ext_temporal_id}[i]$ are equal to hrdDId , hrdQId , and hrdTId , respectively.

When the coded video sequence conforms to one or more of the profiles specified in Annex H and the decoding process specified in Annex H is applied, the following is specified:

- (a) When an access unit contains one or more buffering period SEI messages that are included in MVC scalable nesting SEI messages, the buffering period SEI message that is associated with the operation point being decoded is the buffering period SEI message that initialises the HRD. Let $\text{hrdVId}[i]$ be equal to $\text{sei_op_view_id}[i]$ for all i in the range of 0 to $\text{num_view_components_op_minus1}$, inclusive, and let hrdTId be the value of $\text{sei_op_temporal_id}$, that are associated with the MVC scalable nesting SEI message containing the buffering period SEI message that initialises the HRD.
- (b) The picture timing SEI messages that specify the removal timing of access units from the CPB and output timing from the DPB are the picture timing SEI messages that are included in MVC scalable nesting SEI messages associated with values of $\text{sei_op_view_id}[i]$ equal to $\text{hrdVId}[i]$ for all i in the range of 0 to $\text{num_view_components_op_minus1}$, inclusive, and sei_temporal_id equal to hrdTId .
- (c) The HRD parameters that are used for conformance checking are the HRD parameters included in the MVC video usability information extension of the active MVC sequence parameter set that are associated with values of $\text{vui_mvc_view_id}[i][j]$ for all j in the range of 0 to $\text{vui_mvc_num_target_output_views_minus1}[i]$, inclusive, equal to $\text{hrdVId}[j]$, and the value of $\text{vui_mvc_temporal_id}[i]$ equal to hrdTId . For the specification in this annex, num_units_in_tick , time_scale , $\text{fixed_frame_rate_flag}$, $\text{nal_hrd_parameters_present_flag}$, $\text{vcl_hrd_parameters_present_flag}$, $\text{low_delay_hrd_flag}$, and $\text{pic_struct_present_flag}$ are substituted with the values of $\text{vui_mvc_num_units_in_tick}[i]$, $\text{vui_mvc_time_scale}[i]$, $\text{vui_mvc_fixed_frame_rate_flag}[i]$, $\text{vui_mvc_nal_hrd_parameters_present_flag}[i]$, $\text{vui_mvc_vcl_hrd_parameters_present_flag}[i]$, $\text{vui_mvc_low_delay_hrd_flag}[i]$, and $\text{vui_mvc_pic_struct_present_flag}[i]$, respectively, with i being the value for which $\text{vui_mvc_view_id}[i]$ is equal to $\text{hrdVId}[j]$ for all j in the range of 0 to $\text{vui_mvc_num_target_output_views_minus1}[i]$, inclusive, and $\text{vui_mvc_temporal_id}[i]$ equal to hrdTId .

When the coded video sequence conforms to one or more of the profiles specified in Annex I and the decoding process specified in Annex I is applied, the following is specified:

- (a) When an access unit contains one or more buffering period SEI messages that are included in MVCD scalable nesting SEI messages, the buffering period SEI message that is associated with the operation point being decoded is the buffering period SEI message that initialises the HRD. Let $\text{hrdVId}[i]$ be equal to $\text{sei_op_view_id}[i]$ for all i in the range of 0 to $\text{num_view_components_op_minus1}$, inclusive, and let hrdTId be the value of $\text{sei_op_temporal_id}$, that are associated with the MVCD scalable nesting SEI message containing the buffering period SEI message that initialises the HRD.

- (b) The picture timing SEI messages that specify the removal timing of access units from the CPB and output timing from the DPB are the picture timing SEI messages that are included in MVCD scalable nesting SEI messages associated with values of `sei_op_view_id[i]` equal to `hrdVid[i]` for all *i* in the range of 0 to `num_view_components_op_minus1`, inclusive, and `sei_temporal_id` equal to `hrdTId`.
- (c) The HRD parameter sets that are used for conformance checking are the HRD parameter sets, included in the MVC video usability information extension of the active MVCD sequence parameter set, that are associated with values of `vui_mvc_view_id[i][j]` for all *j* in the range of 0 to `vui_mvc_num_target_output_views_minus1[i]`, inclusive, equal to `hrdVid[j]`, and the value of `vui_mvc_temporal_id[i]` equal to `hrdTId`. For the specification in this annex, `num_units_in_tick`, `time_scale`, `fixed_frame_rate_flag`, `nal_hrd_parameters_present_flag`, `vcl_hrd_parameters_present_flag`, `low_delay_hrd_flag`, and `pic_struct_present_flag` are substituted with the values of `vui_mvc_num_units_in_tick[i]`, `vui_mvc_time_scale[i]`, `vui_mvc_fixed_frame_rate_flag[i]`, `vui_mvc_nal_hrd_parameters_present_flag[i]`, `vui_mvc_vcl_hrd_parameters_present_flag[i]`, `vui_mvc_low_delay_hrd_flag[i]`, and `vui_mvc_pic_struct_present_flag[i]`, respectively, with *i* being the value for which `vui_mvc_view_id[i]` is equal to `hrdVid[j]` for all *j* in the range of 0 to `vui_mvc_num_target_output_views_minus1[i]`, inclusive, and `vui_mvc_temporal_id[i]` equal to `hrdTId`.

The HRD is used to check conformance of bitstreams and decoders as specified in subclauses C.3 and C.4, respectively.

NOTE 9 – While conformance is guaranteed under the assumption that all frame-rates and clocks used to generate the bitstream match exactly the values signalled in the bitstream, in a real system each of these may vary from the signalled or specified value.

All the arithmetic in this annex is done with real values, so that no rounding errors can propagate. For example, the number of bits in a CPB just prior to or after removal of an access unit is not necessarily an integer.

The variable t_c is derived as follows and is called a clock tick:

$$t_c = \text{num_units_in_tick} \div \text{time_scale} \quad (\text{C-1})$$

The following is specified for expressing the constraints in this annex:

- Let access unit *n* be the *n*-th access unit in decoding order with the first access unit being access unit 0.
- Let picture *n* be the primary coded picture or the decoded primary picture of access unit *n*.

C.1 Operation of coded picture buffer (CPB)

The specifications in this subclause apply independently to each set of CPB parameters that is present and to both the Type I and Type II conformance points shown in Figure C-1.

C.1.1 Timing of bitstream arrival

The HRD may be initialised at any one of the buffering period SEI messages. Prior to initialisation, the CPB is empty.

NOTE – After initialisation, the HRD is not initialised again by subsequent buffering period SEI messages.

Each access unit is referred to as access unit *n*, where the number *n* identifies the particular access unit. The access unit that is associated with the buffering period SEI message that initialises the CPB is referred to as access unit 0. The value of *n* is incremented by 1 for each subsequent access unit in decoding order.

The time at which the first bit of access unit *n* begins to enter the CPB is referred to as the initial arrival time $t_{ai}(n)$.

The initial arrival time of access units is derived as follows:

- If the access unit is access unit 0, $t_{ai}(0) = 0$,

– Otherwise (the access unit is access unit n with $n > 0$), the following applies:

- If $\text{cbr_flag}[\text{SchedSelIdx}]$ is equal to 1, the initial arrival time for access unit n , is equal to the final arrival time (which is derived below) of access unit $n - 1$, i.e.,

$$t_{ai}(n) = t_{af}(n - 1) \quad (\text{C-2})$$

- Otherwise ($\text{cbr_flag}[\text{SchedSelIdx}]$ is equal to 0), the initial arrival time for access unit n is derived by

$$t_{ai}(n) = \text{Max}(t_{af}(n - 1), t_{ai, \text{earliest}}(n)) \quad (\text{C-3})$$

where $t_{ai, \text{earliest}}(n)$ is derived as follows:

- If access unit n is not the first access unit of a subsequent buffering period, $t_{ai, \text{earliest}}(n)$ is derived as

$$t_{ai, \text{earliest}}(n) = t_{r,n}(n) - (\text{initial_cpb_removal_delay}[\text{SchedSelIdx}] - \text{initial_cpb_removal_delay_offset}[\text{SchedSelIdx}]) \div 90000 \quad (\text{C-4})$$

with $t_{r,n}(n)$ being the nominal removal time of access unit n from the CPB as specified in subclause C.1.2 and $\text{initial_cpb_removal_delay}[\text{SchedSelIdx}]$ and $\text{initial_cpb_removal_delay_offset}[\text{SchedSelIdx}]$ being specified in the previous buffering period SEI message.

- Otherwise (access unit n is the first access unit of a subsequent buffering period), $t_{ai, \text{earliest}}(n)$ is derived as

$$t_{ai, \text{earliest}}(n) = t_{r,n}(n) - (\text{initial_cpb_removal_delay}[\text{SchedSelIdx}] \div 90000) \quad (\text{C-5})$$

with $\text{initial_cpb_removal_delay}[\text{SchedSelIdx}]$ being specified in the buffering period SEI message associated with access unit n .

The final arrival time for access unit n is derived by

$$t_{af}(n) = t_{ai}(n) + b(n) \div \text{BitRate}[\text{SchedSelIdx}] \quad (\text{C-6})$$

where $b(n)$ is the size in bits of access unit n , counting the bits of the VCL NAL units and the filler data NAL units for the Type I conformance point or all bits of the Type II bitstream for the Type II conformance point, where the Type I and Type II conformance points are as shown in Figure C-1.

The values of SchedSelIdx , $\text{BitRate}[\text{SchedSelIdx}]$, and $\text{CpbSize}[\text{SchedSelIdx}]$ are constrained as follows:

- If the content of the active sequence parameter sets for access unit n and access unit $n - 1$ differ, the HSS selects a value SchedSelIdx1 of SchedSelIdx from among the values of SchedSelIdx provided in the active sequence parameter set for access unit n that results in a $\text{BitRate}[\text{SchedSelIdx1}]$ or $\text{CpbSize}[\text{SchedSelIdx1}]$ for access unit n . The value of $\text{BitRate}[\text{SchedSelIdx1}]$ or $\text{CpbSize}[\text{SchedSelIdx1}]$ may differ from the value of $\text{BitRate}[\text{SchedSelIdx0}]$ or $\text{CpbSize}[\text{SchedSelIdx0}]$ for the value SchedSelIdx0 of SchedSelIdx that was in use for access unit $n - 1$.
- Otherwise, the HSS continues to operate with the previous values of SchedSelIdx , $\text{BitRate}[\text{SchedSelIdx}]$ and $\text{CpbSize}[\text{SchedSelIdx}]$.

When the HSS selects values of $\text{BitRate}[\text{SchedSelIdx}]$ or $\text{CpbSize}[\text{SchedSelIdx}]$ that differ from those of the previous access unit, the following applies:

- the variable $\text{BitRate}[\text{SchedSelIdx}]$ comes into effect at time $t_{ai}(n)$

- the variable $CpbSize[SchedSelIdx]$ comes into effect as follows:
 - If the new value of $CpbSize[SchedSelIdx]$ exceeds the old CPB size, it comes into effect at time $t_{ai}(n)$,
 - Otherwise, the new value of $CpbSize[SchedSelIdx]$ comes into effect at the time $t_r(n)$.

C.1.2 Timing of coded picture removal

When an access unit n is the access unit with n equal to 0 (the access unit that initialises the HRD), the nominal removal time of the access unit from the CPB is specified by

$$t_{r,n}(0) = \text{initial_cpb_removal_delay}[SchedSelIdx] \div 90000 \quad (C-7)$$

When an access unit n is the first access unit of a buffering period that does not initialise the HRD, the nominal removal time of the access unit from the CPB is specified by

$$t_{r,n}(n) = t_{r,n}(n_b) + t_c * \text{cpb_removal_delay}(n) \quad (C-8)$$

where $t_{r,n}(n_b)$ is the nominal removal time of the first access unit of the previous buffering period and $\text{cpb_removal_delay}(n)$ is the value of cpb_removal_delay specified in the picture timing SEI message associated with access unit n .

The nominal removal time $t_{r,n}(n)$ of an access unit n that is not the first access unit of a buffering period is given by

$$t_{r,n}(n) = t_{r,n}(n_b) + t_c * \text{cpb_removal_delay}(n) \quad (C-9)$$

where $t_{r,n}(n_b)$ is the nominal removal time of the first access unit of the current buffering period and $\text{cpb_removal_delay}(n)$ is the value of cpb_removal_delay specified in the picture timing SEI message associated with access unit n .

The removal time of access unit n is specified as follows:

- If $\text{low_delay_hrd_flag}$ is equal to 0 or $t_{r,n}(n) \geq t_{af}(n)$, the removal time of access unit n is specified by

$$t_r(n) = t_{r,n}(n) \quad (C-10)$$

- Otherwise ($\text{low_delay_hrd_flag}$ is equal to 1 and $t_{r,n}(n) < t_{af}(n)$), the removal time of access unit n is specified by

$$t_r(n) = t_{r,n}(n) + t_c * \text{Ceil}((t_{af}(n) - t_{r,n}(n)) \div t_c) \quad (C-11)$$

NOTE – The latter case indicates that the size of access unit n , $b(n)$, is so large that it prevents removal at the nominal removal time.

When an access unit n is the first access unit of a buffering period, n_b is set equal to n at the removal time $t_r(n)$ of the access unit n .

C.2 Operation of the decoded picture buffer (DPB)

The decoded picture buffer contains frame buffers. When a coded video sequence conforming to one or more of the profiles specified in Annex A is decoded by applying the decoding process specified in clauses 2-9, each of the frame buffers may contain a decoded frame, a decoded complementary field pair or a single (non-paired) decoded field that is marked as "used for reference" (reference pictures) or is held for future output (reordered or delayed pictures). When a coded video sequence conforming to one or more of the profiles specified in Annex G is decoded by applying the decoding process specified in Annex G, each frame buffer may contain a decoded frame, a decoded complementary field

pair, a single (non-paired) decoded field, a decoded reference base frame, a decoded reference base complementary field pair or a single (non-paired) decoded reference base field that is marked as "used for reference" (reference pictures) or is held for future output (reordered or delayed pictures). When a coded video sequence conforming to one or more of the profiles specified in Annex H is decoded by applying the decoding process specified in Annex H, each of the frame buffers may contain a decoded frame view component, a decoded complementary field view component pair, or a single (non-paired) decoded field view component that is marked as "used for reference" (reference pictures) or is held for future output (reordered or delayed pictures) or is held as reference for inter-view prediction (inter-view only reference components). When a coded video sequence conforming to one or more of the profiles specified in Annex I is decoded by applying the decoding process specified in Annex I, each of the frame buffers of the texture DPB may contain a, a decoded depth frame view component, a decoded complementary texture field view component pair, or a single (non-paired) decoded texture field view component that is marked as "used for reference" (reference pictures) or is held for future output (reordered or delayed pictures) or is held as reference for inter-view prediction (inter-view only reference components). When a coded video sequence conforming to one or more of the profiles specified in Annex I is decoded by applying the decoding process specified in Annex I, each of the frame buffers of the depth DPB may contain a decoded depth frame view component, a decoded complementary depth field view component pair, or a single (non-paired) decoded depth field view component that is marked as "used for reference" (reference pictures) or is held for future output (reordered or delayed pictures) or is held as reference for inter-view prediction (inter-view only reference components).

Prior to initialisation, the DPB is empty (the DPB fullness is set to zero). The following steps specified in this subclause all happen instantaneously at $t_r(n)$ and in the order listed. When the decoding process specified in Annex H or Annex I is applied, the view components of the current primary coded picture are processed by applying the ordered steps to each view component in increasing order of the associated view order index $VOIdx$. During the invocation of the process for a particular texture view, only the texture view components of the particular view are considered and during the invocation of the process for a particular depth view, only the depth view components of the particular view are considered. For each view component of the current primary coded picture, the corresponding depth view component with the same view order index $VOIdx$, if present, is processed after the texture view component.

1. The process of decoding gaps in `frame_num` and storing "non-existing" frames as specified in subclause C.2.1 is invoked.
2. The picture decoding and output process as specified in subclause C.2.2 is invoked.
3. The process of removing pictures from the DPB before possible insertion of the current picture as specified in subclause C.2.3 is invoked.
4. The process of marking and storing the current decoded picture as specified in subclause C.2.4 is invoked.

NOTE – When the decoding process specified in Annex G is applied, the DPB is only operated for decoded pictures and reference base pictures associated with decoded pictures. The DPB is not operated for layer pictures with `dependency_id` less than `DependencyIdMax` (and associated reference base pictures). All decoded pictures and associated reference base pictures are decoded pictures and associated reference base pictures for `dependency_id` equal to `DependencyIdMax`, which represent the results of the decoding process specified in subclause G.8.

C.2.1 Decoding of gaps in `frame_num` and storage of "non-existing" frames

When the decoding process specified in Annex H is applied, the process specified in this subclause is invoked for a particular view with view order index $VOIdx$, with "picture" being replaced by "view component", "frame" being replaced by "frame view component", and "field" being replaced by "field view component". During the invocation of the process for a particular view, only view components of the particular view are considered and view components of other views are not marked as "unused for reference" or removed from the DPB. When the decoding process specified in Annex I is applied, the process specified in this subclause for Annex H is invoked for particular texture view or depth view with view order index $VOIdx$, with each "view component" being replaced by "texture view component" or "depth view component", "frame view component" being replaced by "texture frame view component" or "depth frame view component", and "field view component" being replaced by "texture field view component". During the invocation of the process for a particular texture view, only the texture view components of the particular view are considered and during the invocation of the process for a particular depth view, only the depth view components of the particular view are considered and view components of other views are not marked as "unused for reference" or removed from the DPB.

The DPB fullness represents the total number of non-empty frame buffers. When the decoding process specified in Annex H is applied, this includes frame buffers that contain view components of other views. When the decoding process specified in Annex I is applied, this includes frame buffers that contain texture or depth view components of other views.

When applicable, gaps in frame_num are detected by the decoding process and the generated frames are marked and inserted into the DPB as specified below.

Gaps in frame_num are detected by the decoding process and the generated frames are marked as specified in subclauses 8.2.5.2 and G.8.2.5.

After the marking of each generated frame, each picture m marked by the "sliding window" process as "unused for reference" is removed from the DPB when it is also marked as "non-existing" or its DPB output time is less than or equal to the CPB removal time of the current picture n ; i.e., $t_{o,dpb}(m) \leq t_r(n)$, or it has OutputFlag equal to 0. When a frame or the last field in a frame buffer is removed from the DPB, the DPB fullness is decremented by one. The "non-existing" generated frame is inserted into the DPB and the DPB fullness is incremented by one.

C.2.2 Picture decoding and output

When the decoding process specified in Annex H is applied, the process specified in this subclause is invoked for a particular view with view order index VOIdx.

When the decoding process specified in Annex I is applied, the process specified in this subclause is invoked for a particular texture view or depth view with view order index VOIdx.

The decoding of the current picture or view component (when applying the decoding process specified in Annex H or Annex I) and the derivation of the DPB output time (if applicable) is specified as follows:

- If the decoding process specified in clause 8 or Annex G is applied, the following applies:
 - The current primary coded picture n is decoded.
 - When picture n has OutputFlag equal to 1, its DPB output time $t_{o,dpb}(n)$ is derived by

$$t_{o,dpb}(n) = t_r(n) + t_c * dpb_output_delay(n) \quad (C-12)$$

where $dpb_output_delay(n)$ is the value of dpb_output_delay specified in the picture timing SEI message associated with access unit n .

- Otherwise (the decoding process specified in Annex H or Annex I is applied), the following applies:
 - The view component with view order index VOIdx of the current primary coded picture n is decoded.
 - When VOIdx is equal to VOIdxMin and any of the view components of picture n has OutputFlag equal to 1, the DPB output time $t_{o,dpb}(n)$ for picture n is derived by Equation C-12, where $dpb_output_delay(n)$ is the value of dpb_output_delay specified in the picture timing SEI message associated with access unit n .

The output of the current picture or view component (when applying the decoding process specified in Annex H or Annex I) is specified as follows:

- If OutputFlag is equal to 1 and $t_{o,dpb}(n) = t_r(n)$, the current picture or view component is output.

NOTE 1 – When the current picture or view component has nal_ref_idc greater than 0 (when using the decoding process specified in Annex G, nal_ref_idc is the syntax element of the target layer representation), it will be stored in the DPB.

- Otherwise, if OutputFlag is equal to 0, the current picture or view component is not output, but it may be stored in the DPB as specified in subclause C.2.4.

- Otherwise (OutputFlag is equal to 1 and $t_{o,dpb}(n) > t_r(n)$), the current picture or view component is output later and will be stored in the DPB (as specified in subclause C.2.4) and is output at time $t_{o,dpb}(n)$ unless indicated not to be output by the decoding or inference of `no_output_of_prior_pics_flag` equal to 1 at a time that precedes $t_{o,dpb}(n)$.

NOTE 2 – When the coded video sequence conforms to a profile specified in Annex H and the decoding process specified in Annex H is used, the view components of all the target output views of a picture are output at the same time instant and in increasing order of the view order index VOIdx.

NOTE 3 – When the coded video sequence conforms to a profile specified in Annex I and the decoding process specified in Annex I is used, the view components of all the target output views of a picture are output at the same time instant and in increasing order of the view order index VOIdx. A depth view component, if present, follows the texture view component within the same view component.

When output, the picture or view component shall be cropped, using the cropping rectangle specified in the active sequence parameter set for the picture or view component.

When the decoding process specified in clause 8 or Annex G is applied, the current picture n is a picture that is output and is not the last picture of the bitstream that is output, the value of $\Delta t_{o,dpb}(n)$ is derived by

$$\Delta t_{o,dpb}(n) = t_{o,dpb}(n_n) - t_{o,dpb}(n) \quad (C-13)$$

where n_n indicates the picture that follows after picture n in output order and has OutputFlag equal to 1.

When the decoding process specified in Annex H or Annex I is applied, the current picture n is a picture that contains at least one view component that is output and the current picture is not the last picture of the bitstream that contains at least one view component that is output and VOIdx is equal to VOIdxMin, the value of $\Delta t_{o,dpb}(n)$ is derived by Equation C-13, where n_n indicates the picture that follows after picture n in output order and contains at least one any view component with OutputFlag equal to 1.

The decoded picture or view component is temporarily stored (not in the DPB).

C.2.3 Removal of pictures from the DPB before possible insertion of the current picture

When the decoding process specified in Annex H is applied, the process specified in this subclause is invoked for a particular view with view order index VOIdx, with "picture" being replaced by "view component", "frame" being replaced by "frame view component", and "field" being replaced by "field view component".

When the decoding process specified in Annex I is applied, the process specified in this subclause for Annex I is invoked for particular texture view and depth view with view order index VOIdx, with each "view component" being replaced by "texture view component" or "depth view component", "frame view component" being replaced by "texture frame view component" or "depth frame view component", and "field view component" being replaced by "texture field view component". During the invocation of the process for a particular texture view, only the texture view components of the particular view are considered and during the invocation of the process for a particular depth view, only the depth view components of the particular view are considered.

When the decoding process specified in Annex H or Annex I is applied, the following process is specified for removing inter-view only reference components of the current access unit from the DPB. By this process, view components of the current view with view order index VOIdx are not removed from the DPB, but inter-view only reference components of other views may be removed. The removal of inter-view only reference components is specified as follows:

- If the view order index VOIdx of the current view is equal to VOIdxMax, all inter-view only reference components m for which any of the following conditions are true are removed from the DPB:
 - OutputFlag is equal to 0,
 - The DPB output time $t_{o,dpb}(m)$ of the picture containing the view component m is less than or equal to the CPB removal time $t_r(n)$ of the current picture.

- Otherwise (the view order index VOIdx of the current view is less than VOIdxMax), all inter-view only reference components m for which both of the following conditions are true are removed from the DPB:
 - OutputFlag is equal to 0 or the DPB output time $t_{o,dpb}(m)$ of the picture containing the view component m is less than or equal to the CPB removal time $t_r(n)$ of the current picture,
 - One of the following conditions is true:
 - The current view component is a view component of an anchor picture and the view_id of the inter-view only reference component m is not equal to any value of anchor_ref_IX[k][j], with X being equal to 0 or 1, k being any integer value greater than the view order index VOIdx of the current view, and j being any integer value in the range of 0 to Max(0, num_anchor_refs_IX[k] – 1), inclusive,
 - The current view component is not a view component of an anchor picture and the view_id of the inter-view only reference component m is not equal to any value of non_anchor_ref_IX[k][j], with X being equal to 0 or 1, k being any integer value greater than the view order index VOIdx of the current view, and j being any integer value in the range of 0 to Max(0, num_non_anchor_refs_IX[k] – 1), inclusive.

When the decoding process specified in Annex H is applied, for the following processes specified in this subclause, only view components of the particular view for which this subclause is invoked are considered, and view components of other views are not marked as "unused for reference" or removed from the DPB. When the decoding process specified in Annex I is applied, for the following processes specified for Annex I in this subclause, during the invocation of the process for a particular texture view, only texture view components of the particular texture view are considered and during the invocation of the process for a particular depth view, only depth view components of the particular depth view are considered, and view components of other views are not marked as "unused for reference" or removed from the DPB. The DPB fullness represents the total number of non-empty frame buffers. When the decoding process specified in Annex H is applied, this includes frame buffers that contain texture view components of other views. When the decoding process specified in Annex I is applied, this includes frame buffers that contain texture or depth view components of other views.

The removal of pictures from the DPB before possible insertion of the current picture proceeds as follows:

- If the decoded picture is an IDR picture the following applies:
 1. All reference pictures in the DPB are marked as "unused for reference" as specified in subclause 8.2.5.1 when a coded video sequence conforming to one or more of the profiles specified in Annex A is decoded by applying the decoding process specified in clauses 2-9, or as specified in subclause G.8.2.4 when a coded video sequence conforming to one or more of the profiles specified in Annex G is decoded by applying the decoding process specified in Annex G, or as specified in subclause H.8.3 when a coded video sequence conforming to one or more of the profiles specified in Annex H is decoded by applying the decoding process specified in Annex H, or as specified in subclause I.8.3 when a coded video sequence conforming to one or more of the profiles specified in Annex I is decoded by applying the decoding process specified in Annex I.
 2. When the IDR picture is not the first IDR picture decoded and the value of PicWidthInMbs or FrameHeightInMbs or max_dec_frame_buffering derived from the active sequence parameter set is different from the value of PicWidthInMbs or FrameHeightInMbs or max_dec_frame_buffering derived from the sequence parameter set that was active for the preceding picture, respectively, no_output_of_prior_pics_flag is inferred to be equal to 1 by the HRD, regardless of the actual value of no_output_of_prior_pics_flag.

NOTE – Decoder implementations should try to handle frame or DPB size changes more gracefully than the HRD in regard to changes in PicWidthInMbs or FrameHeightInMbs.
 3. When no_output_of_prior_pics_flag is equal to 1 or is inferred to be equal to 1, all frame buffers in the DPB are emptied without output of the pictures they contain, and DPB fullness is set to 0.
- Otherwise (the decoded picture is not an IDR picture), the following applies:

- If the slice header of the current picture includes `memory_management_control_operation` equal to 5, all reference pictures in the DPB are marked as "unused for reference".
- Otherwise (the slice header of the current picture does not include `memory_management_control_operation` equal to 5), the decoded reference picture marking process specified in subclause 8.2.5 is invoked when a coded video sequence conforming to one or more of the profiles specified in Annex A is decoded by applying the decoding process specified in clauses 2-9, or the decoded reference picture marking process specified in subclause G.8.2.4 is invoked when a coded video sequence conforming to one or more of the profiles specified in Annex G is decoded by applying the decoding process specified in Annex G, or the decoded reference picture marking process specified in subclause H.8.3 is invoked when a coded video sequence conforming to one or more of the profiles specified in Annex H is decoded by applying the decoding process specified in Annex H, or the decoded reference picture marking process specified in subclause I.8.3 is invoked when a coded video sequence conforming to one or more of the profiles specified in Annex I is decoded by applying the decoding process specified in Annex I.

All pictures m in the DPB, for which all of the following conditions are true, are removed from the DPB:

- picture m is marked as "unused for reference" or picture m is a non-reference picture. When a picture is a reference frame, it is considered to be marked as "unused for reference" only when both of its fields have been marked as "unused for reference",
- picture m is marked as "non-existing" or it has `OutputFlag` equal to 0 or its DPB output time $t_{o,dpb}(m)$ is less than or equal to the CPB removal time $t_r(n)$ of the current picture n .

When a frame or the last field in a frame buffer is removed from the DPB, the DPB fullness is decremented by one.

C.2.4 Current decoded picture marking and storage

When the decoding process specified in Annex H is applied, the process specified in this subclause is invoked for a particular view with view order index $VOIdx$, with "picture" being replaced by "view component", "frame" being replaced by "frame view component", and "field" being replaced by "field view component". When the decoding process specified in Annex I is applied, the process specified in this subclause for Annex I is invoked for particular texture view and depth view with view order index $VOIdx$, with each "view component" being replaced by "texture view component" and "depth view component", "frame view component" being replaced by "texture frame view component" and "depth frame view component", and "field view component" being replaced by "texture field view component". In subclause C.2.4.2, the DPB output time $t_{o,dpb}(n)$ and the CPB removal time $t_r(n)$ of a view component are the DPB output time and the CPB removal time of the picture n containing the view component, depth view component or texture view component.

The marking and storage of the current decoded picture is specified as follows:

- If the current picture is a reference picture, the marking and storage process for reference pictures as specified in subclause C.2.4.1 is invoked.
- Otherwise (the current picture is a non-reference picture), the storage process for non-reference pictures as specified in subclause C.2.4.2 is invoked.

C.2.4.1 Marking and storage of a reference picture into the DPB

The current picture is stored in the DPB as follows:

- If the current decoded picture is a second field (in decoding order) of a complementary reference field pair, and the first field of the pair is still in the DPB, the current decoded picture is stored in the same frame buffer as the first field of the pair.
- Otherwise, the current decoded picture is stored in an empty frame buffer, and the DPB fullness is incremented by one.

When the coded video sequence conforms to one or more of the profiles specified in Annex G and the decoding process specified in Annex G is applied and the current picture has `store_ref_base_pic_flag` equal to 1 (i.e., the current picture is associated with a reference base picture), the associated reference base picture is stored in the DPB as follows:

- If the reference base picture is a second field (in decoding order) of a complementary reference base field pair, and the first field of the pair is still in the DPB, the reference base picture is stored in the same frame buffer as the first field of the pair.
- Otherwise, the reference base picture is stored in an empty frame buffer, and the DPB fullness is incremented by one.

C.2.4.2 Storage of a non-reference picture into the DPB

The variable `storePicFlag` is derived as follows:

- If any of the following conditions are true, `storePicFlag` is set equal to 1:
 - the current picture `n` has `OutputFlag` equal to 1 and $t_{o,dpb}(n) > t_r(n)$,
 - the decoding process specified in Annex H or Annex I is used and the current view component has a view order index `VOIDx` less than `VOIDxMax` and `inter_view_flag` equal to 1.
- Otherwise, `storePicFlag` is set equal to 0.

When `storePicFlag` is equal to 1, the current picture is stored in the DPB as follows:

- If the current decoded picture is a second field (in decoding order) of a complementary non-reference field pair, and the first field of the pair is still in the DPB, the current decoded picture is stored in the same frame buffer as the first field of the pair.
- Otherwise, the current decoded picture is stored in an empty frame buffer, and the DPB fullness is incremented by one.

C.3 Bitstream conformance

A bitstream of coded data conforming to this Recommendation | International Standard fulfils the following requirements.

The bitstream is constructed according to the syntax, semantics, and constraints specified in this Recommendation | International Standard outside of this annex.

The bitstream is tested by the HRD as specified below:

For Type I bitstreams, the number of tests carried out is equal to `cpb_cnt_minus1 + 1` where `cpb_cnt_minus1` is either the syntax element of `hrd_parameters()` following the `vcl_hrd_parameters_present_flag` or is determined by the application by other means not specified in this Recommendation | International Standard. One test is carried out for each bit rate and CPB size combination specified by `hrd_parameters()` following the `vcl_hrd_parameters_present_flag`. Each of these tests is conducted at the Type I conformance point shown in Figure C-1.

For Type II bitstreams there are two sets of tests. The number of tests of the first set is equal to `cpb_cnt_minus1 + 1` where `cpb_cnt_minus1` is either the syntax element of `hrd_parameters()` following the `vcl_hrd_parameters_present_flag` or is determined by the application by other means not specified in this Recommendation | International Standard. One test is carried out for each bit rate and CPB size combination. Each of these tests is conducted at the Type I conformance point shown in Figure C-1. For these tests, only VCL and filler data NAL units are counted for the input bit rate and CPB storage.

The number of tests of the second set, for Type II bitstreams, is equal to `cpb_cnt_minus1 + 1` where `cpb_cnt_minus1` is either the syntax element of `hrd_parameters()` following the `nal_hrd_parameters_present_flag` or is determined by the application by other means not specified in this Recommendation | International Standard. One test is carried out for each

bit rate and CPB size combination specified by `hrd_parameters()` following the `nal_hrd_parameters_present_flag`. Each of these tests is conducted at the Type II conformance point shown in Figure C-1. For these tests, all NAL units (of a Type II NAL unit stream) or all bytes (of a byte stream) are counted for the input bit rate and CPB storage.

NOTE 1 – NAL HRD parameters established by a value of `SchedSelIdx` for the Type II conformance point shown in Figure C-1 are sufficient to also establish VCL HRD conformance for the Type I conformance point shown in Figure C-1 for the same values of `initial_cpb_removal_delay[SchedSelIdx]`, `BitRate[SchedSelIdx]`, and `CpbSize[SchedSelIdx]` for the VBR case (`cbr_flag[SchedSelIdx]` equal to 0). This is because the data flow into the Type I conformance point is a subset of the data flow into the Type II conformance point and because, for the VBR case, the CPB is allowed to become empty and stay empty until the time a next picture is scheduled to begin to arrive. For example, when a coded video sequence conforming to one or more of the profiles specified in Annex A is decoded by applying the decoding process specified in clauses 2-9, when NAL HRD parameters are provided for the Type II conformance point that not only fall within the bounds set for NAL HRD parameters for profile conformance in item j) of subclause A.3.1 or item h) of subclause A.3.3 (depending on the profile in use) but also fall within the bounds set for VCL HRD parameters for profile conformance in item i) of subclause A.3.1 or item g) of subclause A.3.3 (depending on the profile in use), conformance of the VCL HRD for the Type I conformance point is also assured to fall within the bounds of item i) of subclause A.3.1.

For conforming bitstreams, all of the following conditions shall be fulfilled for each of the tests:

1. For each access unit n , with $n > 0$, associated with a buffering period SEI message, with $\Delta t_{g,90}(n)$ specified by

$$\Delta t_{g,90}(n) = 90000 * (t_{r,n}(n) - t_{af}(n-1)) \quad (C-14)$$

the value of `initial_cpb_removal_delay[SchedSelIdx]` shall be constrained as follows:

- If `cbr_flag[SchedSelIdx]` is equal to 0,

$$\text{initial_cpb_removal_delay}[SchedSelIdx] \leq \text{Ceil}(\Delta t_{g,90}(n)) \quad (C-15)$$

- Otherwise (`cbr_flag[SchedSelIdx]` is equal to 1),

$$\text{Floor}(\Delta t_{g,90}(n)) \leq \text{initial_cpb_removal_delay}[SchedSelIdx] \leq \text{Ceil}(\Delta t_{g,90}(n)) \quad (C-16)$$

NOTE 2 – The exact number of bits in the CPB at the removal time of each picture may depend on which buffering period SEI message is selected to initialise the HRD. Encoders must take this into account to ensure that all specified constraints must be obeyed regardless of which buffering period SEI message is selected to initialise the HRD, as the HRD may be initialised at any one of the buffering period SEI messages.

2. A CPB overflow is specified as the condition in which the total number of bits in the CPB is larger than the CPB size. The CPB shall never overflow.
3. A CPB underflow is specified as the condition in which $t_{r,n}(n)$ is less than $t_{af}(n)$. When `low_delay_hrd_flag` is equal to 0, the CPB shall never underflow.
4. The nominal removal times of pictures from the CPB (starting from the second picture in decoding order), shall satisfy the constraints on $t_{r,n}(n)$ and $t_r(n)$ expressed in subclauses A.3.1 through A.3.3 for the profile and level specified in the bitstream when a coded video sequence conforming to one or more of the profiles specified in Annex A is decoded by applying the decoding process specified in clauses 2-9, and they shall satisfy the constraints on $t_{r,n}(n)$ and $t_r(n)$ expressed in subclauses G.10.2.1 and G.10.2.2 for profile and level specified in the bitstream when a coded video sequence conforming to one or more of the profiles specified in Annex G is decoded by applying the decoding process specified in Annex G, and they shall satisfy the constraints on $t_{r,n}(n)$ and $t_r(n)$ expressed in subclause H.10.2 for the profile and level specified in the bitstream when a coded video sequence conforming to one or more of the profiles specified in Annex H is decoded by applying the decoding process specified in Annex H, and they shall satisfy the constraints on $t_{r,n}(n)$ and $t_r(n)$ expressed in subclause I.10.2 for the profile and level specified in the bitstream when a coded video sequence conforming to

one or more of the profiles specified in Annex I is decoded by applying the decoding process specified in Annex I.

5. Immediately after any decoded picture is added to the DPB, the fullness of the DPB shall be less than or equal to the DPB size as constrained by Annexes A, D, and E and subclauses G.10, G.13, G.14, H.10, H.13, H.14, and I.14 for the profile and level specified in the bitstream.
6. All reference pictures shall be present in the DPB when needed for prediction. Each picture shall be present in the DPB at its DPB output time unless it is not stored in the DPB at all, or is removed from the DPB before its output time by one of the processes specified in subclause C.2.
7. The value of $\Delta_{to,dpb}(n)$ as given by Equation C-13, which is the difference between the output time of a picture and that of the first picture following it in output order and having OutputFlag equal to 1, shall satisfy the constraint expressed in subclause A.3.1 for the profile and level specified in the bitstream when a coded video sequence conforming to one or more of the profiles specified in Annex A is decoded by applying the decoding process specified in clauses 2-9, and it shall satisfy the constraint expressed in subclause G.10.2.1 for profile and level specified in the bitstream when a coded video sequence conforming to one or more of the profiles specified in Annex G is decoded by applying the decoding process specified in Annex G, and it shall satisfy the constraints expressed in subclause H.10.2 for the profile and level specified in the bitstream when a coded video sequence conforming to one or more of the profiles specified in Annex H is decoded by applying the decoding process specified in Annex H, and it shall satisfy the constraints expressed in subclause I.10.2 for the profile and level specified in the bitstream when a coded video sequence conforming to one or more of the profiles specified in Annex I is decoded by applying the decoding process specified in Annex I.

C.4 Decoder conformance

A decoder conforming to this Recommendation | International Standard fulfils the following requirements.

A decoder claiming conformance to a specific profile and level shall be able to decode successfully all conforming bitstreams specified for decoder conformance in subclause C.3, provided that all sequence parameter sets and picture parameter sets referred to in the VCL NAL units, and appropriate buffering period and picture timing SEI messages are conveyed to the decoder, in a timely manner, either in the bitstream (by non-VCL NAL units), or by external means not specified by this Recommendation | International Standard.

There are two types of conformance that can be claimed by a decoder: output timing conformance and output order conformance.

To check conformance of a decoder, test bitstreams conforming to the claimed profile and level, as specified in subclause C.3 are delivered by a hypothetical stream scheduler (HSS) both to the HRD and to the decoder under test (DUT). All pictures output by the HRD shall also be output by the DUT and, for each picture output by the HRD, the values of all samples that are output by the DUT for the corresponding picture shall be equal to the values of the samples output by the HRD.

For output timing decoder conformance, the HSS operates as described above, with delivery schedules selected only from the subset of values of SchedSelIdx for which the bit rate and CPB size are restricted as specified in Annex A, Annex G, Annex H and Annex I for the specified profile and level, or with "interpolated" delivery schedules as specified below for which the bit rate and CPB size are restricted as specified in Annex A, Annex G, Annex H, and Annex I. The same delivery schedule is used for both the HRD and DUT.

When the HRD parameters and the buffering period SEI messages are present with cpb_cnt_minus1 greater than 0, the decoder shall be capable of decoding the bitstream as delivered from the HSS operating using an "interpolated" delivery schedule specified as having peak bit rate r , CPB size $c(r)$, and initial CPB removal delay $(f(r) \div r)$ as follows:

$$\alpha = (r - \text{BitRate}[\text{SchedSelIdx} - 1]) \div (\text{BitRate}[\text{SchedSelIdx}] - \text{BitRate}[\text{SchedSelIdx} - 1]), \quad (\text{C-17})$$

$$c(r) = \alpha * \text{CpbSize}[\text{SchedSelIdx}] + (1 - \alpha) * \text{CpbSize}[\text{SchedSelIdx} - 1], \quad (\text{C-18})$$

$$f(r) = \alpha * \text{initial_cpb_removal_delay}[\text{SchedSelIdx}] * \text{BitRate}[\text{SchedSelIdx}] + (1 - \alpha) * \text{initial_cpb_removal_delay}[\text{SchedSelIdx} - 1] * \text{BitRate}[\text{SchedSelIdx} - 1] \quad (\text{C-19})$$

for any $\text{SchedSelIdx} > 0$ and r such that $\text{BitRate}[\text{SchedSelIdx} - 1] \leq r \leq \text{BitRate}[\text{SchedSelIdx}]$ such that r and $c(r)$ are within the limits as specified in Annex A, Annex G, Annex H, and Annex I for the maximum bit rate and buffer size for the specified profile and level.

NOTE 1 – $\text{initial_cpb_removal_delay}[\text{SchedSelIdx}]$ can be different from one buffering period to another and have to be re-calculated.

For output timing decoder conformance, an HRD as described above is used and the timing (relative to the delivery time of the first bit) of picture output is the same for both HRD and the DUT up to a fixed delay.

For output order decoder conformance, the HSS delivers the bitstream to the DUT "by demand" from the DUT, meaning that the HSS delivers bits (in decoding order) only when the DUT requires more bits to proceed with its processing.

NOTE 2 – This means that for this test, the coded picture buffer of the DUT could be as small as the size of the largest access unit.

A modified HRD as described below is used, and the HSS delivers the bitstream to the HRD by one of the schedules specified in the bitstream such that the bit rate and CPB size are restricted as specified in Annex A, Annex G, Annex H, and Annex I. The order of pictures output shall be the same for both HRD and the DUT.

For output order decoder conformance, the HRD CPB size is equal to $\text{CpbSize}[\text{SchedSelIdx}]$ for the selected schedule and the DPB size is equal to MaxDpbFrames . Removal time from the CPB for the HRD is equal to final bit arrival time and decoding is immediate. The operation of the DPB of this HRD is specified in subclause C.4.1.

C.4.1 Operation of the output order DPB

The decoded picture buffer contains frame buffers. When a coded video sequence conforming to one or more of the profiles specified in Annex A is decoded by applying the decoding process specified in clauses 2-9, each of the frame buffers may contain a decoded frame, a decoded complementary field pair or a single (non-paired) decoded field that is marked as "used for reference" or is held for future output (reordered pictures). When a coded video sequence conforming to one or more of the profiles specified in Annex G is decoded by applying the decoding process specified in Annex G, each frame buffer may contain a decoded frame, a decoded complementary field pair, a single (non-paired) decoded field, a decoded reference base frame, a decoded reference base complementary field pair or a single (non-paired) decoded reference base field that is marked as "used for reference" (reference pictures) or is held for future output (reordered or delayed pictures). When a coded video sequence conforming to one or more of the profiles specified in Annex H is decoded by applying the decoding process specified in Annex H, each of the frame buffers may contain a decoded frame view component, a decoded complementary field view component pair, or a single (non-paired) decoded field view component that is marked as "used for reference" (reference pictures) or is held for future output (reordered or delayed pictures) or is held for inter-view prediction (inter-view only reference components). When a coded video sequence conforming to one or more of the profiles specified in Annex I is decoded by applying the decoding process specified in Annex I, each of the frame buffers of the texture DPB may contain a decoded texture frame view component, a decoded complementary texture field view component pair, a single (non-paired) decoded texture field view component that is marked as "used for reference" (reference pictures) or is held for future output (reordered or delayed pictures) or is held for inter-view prediction (inter-view only reference components). When a coded video sequence conforming to one or more of the profiles specified in Annex I is decoded by applying the decoding process specified in Annex I, each of the frame buffers of the depth DPB may contain a decoded depth frame view component, a decoded complementary depth field view component pair, or a single (non-paired) decoded depth field view component that is marked as "used for reference" (reference pictures) or is held for future output (reordered or delayed pictures) or is held as reference for inter-view prediction (inter-view only reference components).

At HRD initialisation, the DPB fullness, measured in non-empty frame buffers, is set equal to 0. The following steps all happen instantaneously when an access unit is removed from the CPB, and in the order listed. When the decoding process specified in Annex H or Annex I is applied, the view components of the current primary coded picture are processed by applying the ordered steps to each view component in increasing order of the associated view order index VOIdx . The invocation of the process for a depth view component, if present, follows the invocation of the process for the texture view component within the same view component.

1. The process of decoding gaps in frame_num and storing "non-existing" frames as specified in subclause C.4.2 is invoked.

2. The picture decoding and output process as specified in subclause C.4.3 is invoked.
3. The process of removing pictures from the DPB before possible insertion of the current picture as specified in subclause C.4.4 is invoked.
4. The process of marking and storing the current decoded picture as specified in subclause C.4.5 is invoked.

NOTE – When the decoding process specified in Annex G is applied, the DPB is only operated for decoded pictures and reference base pictures associated with decoded pictures. The DPB is not operated for layer pictures with dependency_id less than DependencyIdMax (and associated reference base pictures). All decoded pictures and associated reference base pictures are decoded pictures and associated reference base pictures for dependency_id equal to DependencyIdMax, which represent the results of the decoding process specified in subclause G.8.

C.4.2 Decoding of gaps in frame_num and storage of "non-existing" pictures

When the decoding process specified in Annex H is applied, the process specified in this subclause is invoked for a particular view with view order index VOIdx, with "picture" being replaced by "view component", "frame" being replaced by "frame view component", and "field" being replaced by "field view component". During the invocation of the process for a particular view, only view components of the particular view are considered and view components of other views are not marked as "unused for reference" or removed from the DPB.

When the decoding process specified in Annex I is applied, the process specified in this subclause for Annex H is invoked for particular texture view and depth view with view order index VOIdx, with each "view component" being replaced by "texture view component" or "depth view component", "frame view component" being replaced by "texture frame view component" or "depth frame view component", and "field view component" being replaced by "texture field view component". During the invocation of the process for a particular texture view, only the texture view components of the particular view are considered and during the invocation of the process for a particular depth view, only the depth view components of the particular view are considered and view components of other views are not marked as "unused for reference" or removed from the DPB.

The DPB fullness represents the total number of non-empty frame buffers. When the decoding process specified in Annex H is applied, this includes frame buffers that contain view components of other views. When the decoding process specified in Annex I is applied, this includes frame buffers that contain texture or depth view components of other views.

When applicable, gaps in frame_num are detected by the decoding process and the necessary number of "non-existing" frames are inferred in the order specified by the generation of values of UnusedShortTermFrameNum in Equation 7-23 and are marked as specified in subclauses 8.2.5.2 and G.8.2.5. Frame buffers containing a frame or a complementary field pair or a non-paired field which are marked as "not needed for output" and "unused for reference" are emptied (without output), and the DPB fullness is decremented by the number of frame buffers emptied. Each "non-existing" frame is stored in the DPB as follows:

- When there is no empty frame buffer (i.e., DPB fullness is equal to DPB size), the "bumping" process specified in subclause C.4.5.3 is invoked repeatedly until there is an empty frame buffer in which to store the "non-existing" frame.
- The "non-existing" frame is stored in an empty frame buffer and is marked as "not needed for output", and the DPB fullness is incremented by one.

C.4.3 Picture decoding

When the decoding process specified in Annex H is applied, the process specified in this subclause is invoked for a particular view with view order index VOIdx.

When the decoding process specified in Annex I is applied, the process specified for Annex H in this subclause is invoked for a particular texture view and depth view with view order index VOIdx.

The decoding of the current picture or view component (when applying the decoding process specified in Annex H or Annex I) is specified as follows:

- If the decoding process specified in clause 8 or Annex G is applied, the current primary coded picture n is decoded and is temporarily stored (not in the DPB).
- Otherwise (the decoding process specified in Annex H or Annex I is applied), the view component with view order index $VOIdx$ of the current primary coded picture n is decoded and is temporarily stored (not in the DPB).

C.4.4 Removal of pictures from the DPB before possible insertion of the current picture

When the decoding process specified in Annex H is applied, the process specified in this subclause is invoked for a particular view with view order index $VOIdx$, with "picture" being replaced by "view component", "frame" being replaced by "frame view component", and "field" being replaced by "field view component".

When the decoding process specified in Annex I is applied, the process specified in this subclause for Annex H is invoked for particular texture view and depth view with view order index $VOIdx$, with each "view component" being replaced by "texture view component" or "depth view component", "frame view component" being replaced by "texture frame view component" or "depth frame view component", and "field view component" being replaced by "texture field view component". During the invocation of the process for a particular texture view, only the texture view components of the particular view are considered and during the invocation of the process for a particular depth view, only the depth view components of the particular view are considered.

When the decoding process specified in Annex H or Annex I is applied, the following process is specified for emptying frame buffers containing inter-view only reference components of the current access unit. By this process, frame buffers that contain view components of the current view with view order index $VOIdx$ are not emptied, but frame buffers that contain inter-view only reference components of other views may be emptied. The process is specified as follows:

- If the view order index $VOIdx$ of the current view is equal to $VOIdxMax$, all frame buffers containing a frame or a complementary field pair or a non-paired field which are marked as "not needed for output" and "unused for reference" are emptied (without output), and the DPB fullness is decremented by the number of frame buffers emptied.

NOTE 1 – At this stage of the process, all frame buffers that contain a frame or a complementary field pair or a non-paired field marked as "not needed for output" and "unused for reference" are frame buffers that contain an inter-view only reference component (of the current access unit and a view with view order index less than $VOIdx$) with $OutputFlag$ equal to 0.

- Otherwise (the view order index $VOIdx$ of the current view is less than $VOIdxMax$), frame buffers containing a frame or a complementary field pair or a non-paired field for which both of the following conditions are true are emptied (without output), and the DPB fullness is decremented by the number of frame buffers emptied:
 - the frame or complementary field pair or non-paired field is marked as "not needed for output" and "unused for reference",

NOTE 2 – At this stage of the process, all frame buffers that contain a frame or a complementary field pair or a non-paired field marked as "not needed for output" and "unused for reference" are frame buffers that contain an inter-view only reference component (of the current access unit and a view with view order index less than $VOIdx$) with $OutputFlag$ equal to 0.

- one of the following conditions is true:
 - the current view component is a view component of an anchor picture and the $view_id$ of the frame or complementary field pair or non-paired field is not equal to any value of $anchor_ref_IX[k][j]$, with X being equal to 0 or 1, k being any integer value greater than the view order index $VOIdx$ of the current view, and j being any integer value in the range of 0 to $Max(0, num_anchor_refs_IX[k] - 1)$, inclusive,

- the current view component is not a view component of an anchor picture and the view_id of the frame or complementary field pair or non-paired field is not equal to any value of non_anchor_ref_IX[k][j], with X being equal to 0 or 1, k being any integer value greater than the view order index VOIdx of the current view, and j being any integer value in the range of 0 to Max(0, num_non_anchor_refs_IX[k] – 1), inclusive.

When the decoding process specified in Annex H or Annex I is applied, for the following processes specified in this subclause, only view components of the particular view for which this subclause is invoked are considered, and frame buffers containing view components of other views are not emptied. The DPB fullness represents the total number of non-empty frame buffers, including frame buffers that contain view components of other views.

The removal of pictures from the DPB before possible insertion of the current picture proceeds as follows:

- If the decoded picture is an IDR picture the following applies:
 1. All reference pictures in the DPB are marked as "unused for reference" as specified in subclause 8.2.5 when a coded video sequence conforming to one or more of the profiles specified in Annex A is decoded by applying the decoding process specified in clauses 2-9, or as specified in subclause G.8.2.4 when a coded video sequence conforming to one or more of the profiles specified in Annex G is decoded by applying the decoding process specified in Annex G, or as specified in subclause H.8.3 when a coded video sequence conforming to one or more of the profiles specified in Annex H is decoded by applying the decoding process specified in Annex H, or as specified in subclause I.8.3 when a coded video sequence conforming to one or more of the profiles specified in Annex I is decoded by applying the decoding process specified in Annex I.
 2. When the IDR picture is not the first IDR picture decoded and the value of PicWidthInMbs or FrameHeightInMbs or max_dec_frame_buffering derived from the active sequence parameter set is different from the value of PicWidthInMbs or FrameHeightInMbs or max_dec_frame_buffering derived from the sequence parameter set that was active for the preceding picture, respectively, no_output_of_prior_pics_flag is inferred to be equal to 1 by the HRD, regardless of the actual value of no_output_of_prior_pics_flag.

NOTE 3 – Decoder implementations should try to handle changes in the value of PicWidthInMbs or FrameHeightInMbs or max_dec_frame_buffering more gracefully than the HRD.

 3. When no_output_of_prior_pics_flag is equal to 1 or is inferred to be equal to 1, all frame buffers in the DPB are emptied without output of the pictures they contain, and DPB fullness is set to 0.
- Otherwise (the decoded picture is not an IDR picture), the decoded reference picture marking process is invoked as specified in subclause 8.2.5 when a coded video sequence conforming to one or more of the profiles specified in Annex A is decoded by applying the decoding process specified in clauses 2-9, or as specified in subclause G.8.2.4 when a coded video sequence conforming to one or more of the profiles specified in Annex G is decoded by applying the decoding process specified in Annex G, or as specified in subclause H.8.3 when a coded video sequence conforming to one or more of the profiles specified in Annex H is decoded by applying the decoding process specified in Annex H, or as specified in subclause I.8.3 when a coded video sequence conforming to one or more of the profiles specified in Annex I is decoded by applying the decoding process specified in Annex I. Frame buffers containing a frame or a complementary field pair or a non-paired field which are marked as "not needed for output" and "unused for reference" are emptied (without output), and the DPB fullness is decremented by the number of frame buffers emptied.

When the current picture has a memory_management_control_operation equal to 5 or is an IDR picture for which no_output_of_prior_pics_flag is not equal to 1 and is not inferred to be equal to 1, the following two steps are performed.

1. Frame buffers containing a frame or a complementary field pair or a non-paired field which are marked as "not needed for output" and "unused for reference" are emptied (without output), and the DPB fullness is decremented by the number of frame buffers emptied.
2. All non-empty frame buffers in the DPB are emptied by repeatedly invoking the "bumping" process specified in subclause C.4.5.3, and the DPB fullness is set to 0.

C.4.5 Current decoded picture marking and storage

When the decoding process specified in Annex H is applied, the process specified in this subclause is invoked for a particular view with view order index VOIdx, with "picture" being replaced by "view component", "frame" being replaced by "frame view component", and "field" being replaced by "field view component". During the invocation of the process for a particular view, only view components of the particular view are considered and frame buffers containing view components of other views are not emptied.

When the decoding process specified in Annex I is applied, the process specified in this subclause for Annex H is invoked for particular texture view and depth view with view order index VOIdx, with each "view component" being replaced by "texture view component" or "depth view component", "frame view component" being replaced by "texture frame view component" or "depth frame view component", and "field view component" being replaced by "texture field view component". During the invocation of the process for a particular texture view, only the texture view components of the particular view are considered and during the invocation of the process for a particular depth view, only the depth view components of the particular view are considered and frame buffers containing view components of other views are not emptied.

The DPB fullness represents the total number of non-empty frame buffers. When the decoding process specified in Annex H is applied, this includes frame buffers that contain view components of other views. When the decoding process specified in Annex I is applied, this includes frame buffers that contain texture or depth view components of other views.

The marking and storage of the current decoded picture is specified as follows:

- If the current picture is a reference picture, the storage and marking process for decoded reference pictures as specified in subclause C.4.5.1 is invoked.
- Otherwise (the current picture is a non-reference picture), the storage and marking process for decoded non-reference pictures as specified in subclause C.4.5.2 is invoked.

C.4.5.1 Storage and marking of a reference decoded picture into the DPB

The current picture is stored in the DPB as follows:

- If the current decoded picture is the second field (in decoding order) of a complementary reference field pair, and the first field of the pair is still in the DPB, the current picture is stored in the same frame buffer as the first field of the pair and the following applies:
 - If the current decoded picture has OutputFlag equal to 1, it is marked as "needed for output".
 - Otherwise (the current decoded picture has OutputFlag equal to 0), it is marked as "not needed for output".
- Otherwise, the following operations are performed:
 1. When there is no empty frame buffer (i.e., DPB fullness is equal to DPB size), the "bumping" process specified in subclause C.4.5.3 is invoked repeatedly until there is an empty frame buffer in which to store the current decoded picture.
 2. The current decoded picture is stored in an empty frame buffer, the DPB fullness is incremented by one, and the following applies:
 - If the current decoded picture has OutputFlag equal to 1, it is marked as "needed for output".
 - Otherwise (the current decoded picture has OutputFlag equal to 0), it is marked as "not needed for output".

When the coded video sequence conforms to one or more of the profiles specified in Annex G and the decoding process specified in Annex G is applied and the current picture has `store_ref_base_pic_flag` equal to 1 (i.e., the current picture is associated with a reference base picture), the associated reference base picture is stored in the DPB as follows:

- If the reference base picture is a second field (in decoding order) of a complementary reference base field pair, and the first field of the pair is still in the DPB, the reference base picture is stored in the same frame buffer as the first field of the pair and marked as "not needed for output".
- Otherwise, the following operations are performed:
 1. When there is no empty frame buffer (i.e., DPB fullness is equal to DPB size), the "bumping" process specified in subclause C.4.5.3 is invoked repeatedly until there is an empty frame buffer in which to store the reference base picture.
 2. The reference base picture is stored in an empty frame buffer and marked as "not needed for output" and the DPB fullness is incremented by one.

C.4.5.2 Storage and marking of a non-reference decoded picture into the DPB

The current picture is associated with a variable `StoreInterViewOnlyRefFlag`, which is derived as follows:

- If the decoding process specified in Annex H or Annex I is applied, the current view component has a view order index `VOIdx` less than `VOIdxMax` and `inter_view_flag` equal to 1, `StoreInterViewOnlyRefFlag` is set equal to 1.
- Otherwise, `StoreInterViewOnlyRefFlag` is set equal to 0.

The current picture is stored in the DPB or output as follows:

- If the current decoded picture is the second field (in decoding order) of a complementary non-reference field pair and the first field of the pair is still in the DPB, the current picture is stored in the same frame buffer as the first field of the pair and the following applies:
 - If the current decoded picture has `OutputFlag` equal to 1, it is marked as "needed for output".
 - Otherwise (the current decoded picture has `OutputFlag` equal to 0), it is marked as "not needed for output".
- Otherwise, if the current picture has `OutputFlag` equal to 0 and `StoreInterViewOnlyRefFlag` equal to 0, the DPB is not modified and the current picture is not output.
- Otherwise, if the current picture has `StoreInterViewOnlyRefFlag` equal to 1, the following operations are performed:
 1. When there is no empty frame buffer (i.e., DPB fullness is equal to DPB size), the "bumping" process specified in subclause C.4.5.3 is invoked repeatedly until there is an empty frame buffer in which to store the current decoded picture.
 2. The current decoded picture is stored in an empty frame buffer, the DPB fullness is incremented by one, and the following applies:
 - If the current decoded picture has `OutputFlag` equal to 1, it is marked as "needed for output".
 - Otherwise (the current decoded picture has `OutputFlag` equal to 0), it is marked as "not needed for output".
- Otherwise, the following operations are performed repeatedly until the current decoded picture has been cropped and output or has been stored in the DPB:
 - If there is no empty frame buffer (i.e., DPB fullness is equal to DPB size), the following applies:

- If the current picture does not have a lower value of `PicOrderCnt()` than all pictures in the DPB that are marked as "needed for output", the "bumping" process described in subclause C.4.5.3 is performed.
- Otherwise (the current picture has a lower value of `PicOrderCnt()` than all pictures in the DPB that are marked as "needed for output"), the current picture is cropped, using the cropping rectangle specified in the active sequence parameter set for the picture and the cropped picture is output.
- Otherwise (there is an empty frame buffer, i.e., DPB fullness is less than DPB size), the current decoded picture is stored in an empty frame buffer and is marked as "needed for output", and the DPB fullness is incremented by one.

C.4.5.3 "Bumping" process

When the decoding process specified in Annex H is applied, the process specified in this subclause is invoked for a particular view with view order index `VOIdx`, with "picture" being replaced by "view component", "frame" being replaced by "frame view component", and "field" being replaced by "field view component". During the invocation of the process for a particular view, only view components of the particular view are considered and frame buffers containing view components of other views are not emptied.

When the decoding process specified in Annex I is applied, the process specified in this subclause for Annex H is invoked for particular texture view and depth view with view order index `VOIdx`, with each "view component" being replaced by "texture view component" or "depth view component", "frame view component" being replaced by "texture frame view component" or "depth frame view component", and "field view component" being replaced by "texture field view component". During the invocation of the process for a particular texture view, only the texture view components of the particular view are considered while respective depth view components may be cropped and output too. During the invocation of the process for a particular depth view, only the depth view components of the particular view are considered and frame buffers containing view components of other views are not emptied. The DPB fullness represents the total number of non-empty frame buffers, including frame buffers that contain view components of other views, for the texture DPB or the depth DPB depending on whether the process is invoked for a texture view or a depth view, respectively.

The DPB fullness represents the total number of non-empty frame buffers. When the decoding process specified in Annex H is applied, this includes frame buffers that contain view components of other views. When the decoding process specified in Annex I is applied, this includes frame buffers that contain texture or depth view components of other views.

The "bumping" process is invoked in the following cases.

- There is no empty frame buffer (i.e., DPB fullness is equal to DPB size) and an empty frame buffer is needed for storage of an inferred "non-existing" frame, as specified in subclause C.4.2.
- The current picture is an IDR picture and `no_output_of_prior_pics_flag` is not equal to 1 and is not inferred to be equal to 1, as specified in subclause C.4.4.
- The current picture has `memory_management_control_operation` equal to 5, as specified in subclause C.4.4.
- There is no empty frame buffer (i.e., DPB fullness is equal to DPB size) and an empty frame buffer is needed for storage of a decoded (non-IDR) reference picture or a reference base picture, as specified in subclause C.4.5.1.
- There is no empty frame buffer (i.e., DPB fullness is equal to DPB size) and the current picture is a non-reference picture that is not the second field of a complementary non-reference field pair and the current picture has `OutputFlag` equal to 1 and there are pictures in the DPB that are marked as "needed for output" that precede the current non-reference picture in output order, as specified in subclause C.4.5.2, so an empty buffer is needed for storage of the current picture.

- There is no empty frame buffer (i.e., DPB fullness is equal to DPB size) and the current picture is a non-reference picture that is not the second field of a complementary non-reference field pair and the current picture has StoreInterViewOnlyRefFlag equal to 1, as specified in subclause C.4.5.2, so an empty buffer is needed for storage of the current picture.

The "bumping" process consists of the following ordered steps:

1. The picture or complementary reference field pair that is considered first for output is selected as follows:
 - a. The frame buffer is selected that contains the picture having the smallest value of PicOrderCnt() of all pictures in the DPB marked as "needed for output".
 - b. Depending on the frame buffer, the following applies:
 - If this frame buffer contains a complementary non-reference field pair with both fields marked as "needed for output" and both fields have the same PicOrderCnt(), the first of these two fields in decoding order is considered first for output.
 - Otherwise, if this frame buffer contains a complementary reference field pair with both fields marked as "needed for output" and both fields have the same PicOrderCnt(), the entire complementary reference field pair is considered first for output.

NOTE – When the two fields of a complementary reference field pair have the same value of PicOrderCnt(), this "bumping" process will output these pictures together, although the two fields have different output times from a decoder that satisfies output timing conformance criteria (as specified in subclause C.2.2).

 - Otherwise, the picture in this frame buffer that has the smallest value of PicOrderCnt() is considered first for output.
2. Depending on whether a single picture or a complementary reference field pair is considered for output, the following applies:
 - If a single picture is considered first for output, this picture is cropped, using the cropping rectangle specified in the active sequence parameter set for the picture, the cropped picture is output, and the picture is marked as "not needed for output".
 - Otherwise (a complementary reference field pair is considered first for output), the two fields of the complementary reference field pair are both cropped, using the cropping rectangle specified in the active sequence parameter set for the pictures, the two fields of the complementary reference field pair are output together, and both fields of the complementary reference field pair are marked as "not needed for output".
3. When there is a single depth view component or a complementary depth view component pair having the same values of view_id and PicOrderCnt() as the single picture or complementary reference field pair considered for output, the single depth view component or complementary depth view component pair are output as in step 2.
4. The frame buffer that included the picture or complementary reference field pair that was cropped and output is checked, and when any of the following conditions are true, the frame buffer is emptied and the DPB fullness is decremented by 1:
 - The frame buffer contains a non-reference non-paired field.
 - The frame buffer contains a non-reference frame.
 - The frame buffer contains a complementary non-reference field pair with both fields marked as "not needed for output".
 - The frame buffer contains a non-paired reference field marked as "unused for reference".
 - The frame buffer contains a reference frame with both fields marked as "unused for reference".

- In Table D.1, add the following after "display_orientation (payloadSize)":*

else if(payloadType == 48)		
mvcd_scalable_nesting(payloadSize) /* specified in Annex I */	5	
else if (payloadType == 49)		
mvcd_view_scalability_info(payloadSize) /* specified in Annex I */	5	
else if (payloadType == 50)		
depth_representation_info(payloadSize) /* specified in Annex I */	5	
else if (payloadType == 51)		
three_dimensional_reference_displays_info(payloadSize) /* specified in Annex I */	5	
else if (payloadType == 52)		
depth_timing(payloadSize) /* specified in Annex I */	5	
else if(payloadType == 53)		
depth_sampling_info(payloadSize) /* specified in Annex I */	5	

Replace the following sentence

with

Add the following paragraph after the paragraph containing the sentence above:

The multiview acquisition information SEI message may be nested in an MVCD scalable nesting SEI message to indicate parameters of the acquisition environment of texture and depth views. When present as a nested SEI message, the multiview acquisition information SEI message is recommended be associated with an IDR access unit and may be associated with any access unit. When present as a nested SEI message, the information indicated in the SEI message applies from the access unit associated with the SEI message to the next access unit, in decoding order, containing an SEI message of the same type, exclusive, or to the end of the coded video sequence, whichever is earlier in decoding order.

Replace the two paragraphs specifying the semantics of "num_views_minus1" with the following:

num_views_minus1 shall be equal to the value of the syntax element num_views_minus1 in the active MVC sequence parameter set for the coded video sequence when the SEI message is not nested. When the SEI message is nested in an MVCD scalable nesting SEI message, num_views_minus1 shall be equal to the value of num_view_components_minus1 of the containing MVCD scalable nesting SEI message. The value of num_views_minus1 shall be in the range of 0 to 1023, inclusive.

When the SEI message is not nested, the loop index *i* in the subsequent syntax elements indicates the view order index derived from the active MVC sequence parameter set. When the SEI message is nested in an MVCD scalable nesting SEI message, the loop index *i* in the subsequent syntax elements indicates the view with view_id equal to sei_view_id[*i*] of the containing MVCD scalable nesting SEI message.

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC 14496-10:2012/Amd.2:2013

Add Annex I:

Annex I

Multiview and depth video coding

(This annex forms an integral part of this Recommendation | International Standard.)

This annex specifies multiview video coding with depth information, referred to as MVCD.

I.1 Scope

Bitstreams and decoders conforming to the profile specified in this annex are completely specified in this annex with reference made to clauses 2-9 and Annexes A-H.

I.2 Normative references

The specifications in clause 2 apply.

I.3 Definitions

For the purpose of this annex, the following definitions apply in addition to the definitions in subclause H.3. These definitions are either not present in subclause H.3 or replace definitions in subclause H.3.

I.3.1 depth field view: A *depth view component* of a *field*.

I.3.2 depth frame view: A *depth view component* of a *frame*.

I.3.3 depth view: A sequence of *depth view components* associated with an identical value of *view_id*.

I.3.4 depth view component: A *coded representation* of the depth of a view in a single *access unit*.

I.3.5 inter-view only reference component: A *view component*, *texture view component*, or *depth view component* coded with *nal_ref_idc* equal to 0 and *inter_view_flag* equal to 1. An *inter-view only reference component* contains samples that may be used for *inter-view prediction* in the *decoding process* of subsequent *view components* in *decoding order*, but are not used for *inter prediction* by any *view components*. *Inter-view only reference components* are *non-reference pictures*.

I.3.6 inter-view reference component: A *view component*, *texture view component*, or *depth view component* coded with *nal_ref_idc* greater than 0 and *inter_view_flag* equal to 1. An *inter-view reference component* contains samples that may be used for *inter prediction* of subsequent *pictures* in *decoding order* and *inter-view prediction* of subsequent *view components* in *decoding order*. *Inter-view reference components* are *reference pictures*.

I.3.7 MVCD operation point: An operation point for which each target output view includes a *texture view* or a *depth view* or both a *texture view* and a *depth view*.

I.3.8 MVCD sequence parameter set: A collective term for *sequence parameter set* or *subset sequence parameter set*.

I.3.9 MVCD sequence parameter set RBSP: A collective term for *sequence parameter set RBSP* or *subset sequence parameter set RBSP*.

I.3.10 reference picture: A *view component*, *texture view component*, or *depth view component* coded with *nal_ref_idc* greater than 0. A *reference picture* contains samples that may be used for *inter prediction* in the *decoding process* of subsequent *view components* in *decoding order*. A *reference picture* may be an *inter-view reference component*, in which case the samples contained in the *reference picture* may also be used for *inter-view prediction* in the *decoding process* of subsequent *view components* in *decoding order*.

- I.3.11 stereoscopic texture bitstream:** A *bitstream* containing two *texture views* and conforming to one of the *profiles* specified in Annex H.
- I.3.12 texture field view component:** A *texture view component* of a *field*.
- I.3.13 texture frame view component:** A *texture view component* of a *frame*.
- I.3.14 texture view:** A sequence of *texture view components* associated with an identical value of *view_id*.
- I.3.15 texture view component:** A *coded representation* of the texture of a view in a single *access unit*.
- I.3.16 view:** A *texture view* and a *depth view* with the same value of *view_id*, unless explicitly limited to either *texture view* or *depth view*.
- I.3.17 view component:** A *coded representation* of a view in a single *access unit*. A *view component* may consist of a *texture view component* and a *depth view component*.
- I.3.18 view component pair:** A *texture view component* and a *depth view component* of the same view within the same *access unit*.

I.4 Abbreviations

The specifications in clause 4 apply.

I.5 Conventions

The specifications in clause 5 apply.

I.6 Source, coded, decoded and output data formats, scanning processes, and neighbouring relationships

The specifications in clause 6 apply with substitution of MVCD sequence parameter set for sequence parameter set.

I.7 Syntax and semantics

This clause specifies syntax and semantics for coded video sequences that conform to one or more of the profiles specified in this annex.

I.7.1 Method of specifying syntax in tabular form

The specifications in subclause H.7.1 apply.

I.7.2 Specification of syntax functions, categories, and descriptors

The specifications in subclause H.7.2 apply.

I.7.3 Syntax in tabular form

I.7.3.1 NAL unit syntax

The syntax table is specified in subclause H.7.3.1.

I.7.3.1.1 NAL unit header MVC extension syntax

The syntax table is specified in subclause H.7.3.1.1.

I.7.3.2 Raw byte sequence payloads and RBSP trailing bits syntax**I.7.3.2.1 Sequence parameter set RBSP syntax**

The syntax table is specified in subclause H.7.3.2.1.

I.7.3.2.1.1 Sequence parameter set data syntax

The syntax table is specified in subclause H.7.3.2.1.1.

I.7.3.2.1.1.1 Scaling list syntax

The syntax table is specified in subclause H.7.3.2.1.1.1.

I.7.3.2.1.2 Sequence parameter set extension RBSP syntax

The syntax table is specified in subclause H.7.3.2.1.2.

I.7.3.2.1.3 Subset sequence parameter set RBSP syntax

The syntax table is specified in subclause H.7.3.2.1.3.

I.7.3.2.1.4 Sequence parameter set MVC extension syntax

The syntax table is specified in subclause H.7.3.2.1.4.

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC 14496-10:2012/Amd 2:2013

I.7.3.2.1.5 Sequence parameter set MVCD extension syntax

seq_parameter_set_mvcd_extension() {	C	Descriptor
num_views_minus1	0	ue(v)
for(i = 0; NumDepthViews = 0; i <= num_views_minus1; i++) {		
view_id[i]	0	ue(v)
depth_view_present_flag[i]	0	u(1)
DepthViewId[NumDepthViews] = view_id[i]		
NumDepthViews += depth_view_present_flag[i]		
texture_view_present_flag[i]	0	u(1)
}		
for(i = 1; i <= num_views_minus1; i++)		
if(depth_view_present_flag[i]) {		
num_anchor_refs_l0[i]	0	ue(v)
for(j = 0; j < num_anchor_refs_l0[i]; j++)		
anchor_ref_l0[i][j]	0	ue(v)
num_anchor_refs_l1[i]	0	ue(v)
for(j = 0; j < num_anchor_refs_l1[i]; j++)		
anchor_ref_l1[i][j]	0	ue(v)
}		
for(i = 1; i <= num_views_minus1; i++)		
if(depth_view_present_flag[i]) {		
num_non_anchor_refs_l0[i]	0	ue(v)
for(j = 0; j < num_non_anchor_refs_l0[i]; j++)		
non_anchor_ref_l0[i][j]	0	ue(v)
num_non_anchor_refs_l1[i]	0	ue(v)
for(j = 0; j < num_non_anchor_refs_l1[i]; j++)		
non_anchor_ref_l1[i][j]	0	ue(v)
}		
num_level_values_signalled_minus1	0	ue(v)
for(i = 0; i <= num_level_values_signalled_minus1; i++) {		
level_idc[i]	0	u(8)
num_applicable_ops_minus1[i]	0	ue(v)
for(j = 0; j <= num_applicable_ops_minus1[i]; j++) {		
applicable_op_temporal_id[i][j]	0	u(3)
applicable_op_num_target_views_minus1[i][j]	0	ue(v)
for(k = 0; k <= applicable_op_num_target_views_minus1[i][j]; k++) {		
applicable_op_target_view_id[i][j][k]	0	ue(v)
applicable_op_depth_flag[i][j][k]	0	u(1)
applicable_op_texture_flag[i][j][k]	0	u(1)
}		
applicable_op_num_texture_views_minus1[i][j]	0	ue(v)
applicable_op_num_depth_views[i][j]	0	ue(v)
}		
}		
mvcd_vui_parameters_present_flag	0	u(1)
if(mvcd_vui_parameters_present_flag == 1)		
mvcd_vui_parameters_extension()		

texture_vui_parameters_present_flag	0	u(1)
if(texture_vui_parameters_present_flag == 1)		
mvc_vui_parameters_extension()	0	
}		

I.7.3.2.2 Picture parameter set RBSP syntax

The syntax table is specified in subclause H.7.3.2.2.

I.7.3.2.3 Supplemental enhancement information RBSP syntax

The syntax table is specified in subclause H.7.3.2.3.

I.7.3.2.3.1 Supplemental enhancement information message syntax

The syntax table is specified in subclause H.7.3.2.3.1.

I.7.3.2.4 Access unit delimiter RBSP syntax

The syntax table is specified in subclause H.7.3.2.4.

I.7.3.2.5 End of sequence RBSP syntax

The syntax table is specified in subclause H.7.3.2.5.

I.7.3.2.6 End of stream RBSP syntax

The syntax table is specified in subclause H.7.3.2.6.

I.7.3.2.7 Filler data RBSP syntax

The syntax table is specified in subclause H.7.3.2.7.

I.7.3.2.8 Slice layer without partitioning RBSP syntax

The syntax table is specified in subclause H.7.3.2.8.

I.7.3.2.9 Slice data partition RBSP syntax

Slice data partition syntax is not present in coded video sequences conforming to one or more of the profiles specified in this annex.

I.7.3.2.10 RBSP slice trailing bits syntax

The syntax table is specified in subclause H.7.3.2.10.

I.7.3.2.11 RBSP trailing bits syntax

The syntax table is specified in subclause H.7.3.2.11.

I.7.3.2.12 Prefix NAL unit RBSP syntax

The syntax table is specified in subclause H.7.3.2.12.

I.7.3.2.13 Slice layer extension RBSP syntax

The syntax table is specified in subclause H.7.3.2.13.

I.7.3.3 Slice header syntax

The syntax table is specified in subclause H.7.3.3.

I.7.3.3.1 Reference picture list modification syntax

The syntax table is specified in subclause H.7.3.3.1.

I.7.3.3.1.1 Reference picture list MVC modification syntax

The syntax table is specified in subclause H.7.3.3.1.1

I.7.3.3.2 Prediction weight table syntax

The syntax table is specified in subclause H.7.3.3.2.

I.7.3.3.3 Decoded reference picture marking syntax

The syntax table is specified in subclause H.7.3.3.3.

I.7.3.4 Slice data syntax

The syntax table is specified in subclause H.7.3.4.

I.7.3.5 Macroblock layer syntax

The syntax table is specified in subclause H.7.3.5.

I.7.3.5.1 Macroblock prediction syntax

The syntax table is specified in subclause H.7.3.5.1.

I.7.3.5.2 Sub-macroblock prediction syntax

The syntax table is specified in subclause H.7.3.5.2.

I.7.3.5.3 Residual data syntax

The syntax table is specified in subclause H.7.3.5.3.

I.7.3.5.3.1 Residual luma syntax

The syntax table is specified in subclause H.7.3.5.3.1.

I.7.3.5.3.2 Residual block CAVLC syntax

The syntax table is specified in subclause H.7.3.5.3.2.

I.7.3.5.3.3 Residual block CABAC syntax

The syntax table is specified in subclause H.7.3.5.3.3.

I.7.4 Semantics

Semantics associated with the syntax structures and syntax elements within these structures (in subclause I.7.3 and in subclause H.7.3 by reference in subclause I.7.3) are specified in this subclause and by reference to subclause I.7.4. When the semantics of a syntax element are specified using a table or a set of tables, any values that are not specified in the table(s) shall not be present in the bitstream unless otherwise specified in this Recommendation | International Standard.

I.7.4.1 NAL unit semantics

The semantics for the syntax elements in subclause I.7.3.1 are specified in subclause H.7.3.1.

I.7.4.1.1 NAL unit header MVC extension semantics

The semantics for the syntax elements in subclause I.7.3.1.1 are specified in subclause H.7.3.1.1.

I.7.4.1.2 Order of NAL units and association to coded pictures, access units, and video sequences

This subclause specifies constraints on the order of NAL units in the bitstream. Any order of NAL units in the bitstream obeying these constraints is referred to in the text as the decoding order of NAL units. Within a NAL unit, the syntax in subclauses 7.3, D.1, E.1, H.7.3, H.13.1, H.14.1, I.13.1 and I.14.1 specifies the decoding order of syntax elements. Decoders shall be capable of receiving NAL units and their syntax elements in decoding order.

I.7.4.1.2.1 Order of MVCD sequence parameter set RBSPs and picture parameter set RBSPs and their activation

NOTE 1 – The sequence and picture parameter set mechanism decouples the transmission of infrequently changing information from the transmission of coded macroblock data. Sequence and picture parameter sets may, in some applications, be conveyed "out-of-band" using a reliable transport mechanism.

A picture parameter set RBSP includes parameters that can be referred to by the coded slice NAL units of one or more texture view or depth view components of one or more coded pictures.

Each picture parameter set RBSP is initially considered not active at the start of the operation of the decoding process. At most one picture parameter set RBSP is considered as the active picture parameter set RBSP at any given moment during the operation of the decoding process, and when any particular picture parameter set RBSP becomes the active picture parameter set RBSP, the previously-active picture parameter set RBSP (if any) is deactivated.

In addition to the active picture parameter set RBSP, zero or more picture parameter set RBSPs may be specifically active for texture view components (with a particular value of VOIdx less than or equal to VOIdxMax) that belong to the target output views or that may be referred to through inter-view prediction in decoding texture view components belonging to the target output views. Such a picture parameter set RBSP is referred to as the active texture picture parameter set RBSP for the particular value of VOIdx. The restrictions on active picture parameter set RBSPs also apply to active texture picture parameter set RBSPs for a particular value of VOIdx.

Furthermore, zero or more picture parameter set RBSPs may be specifically active for depth view components (with a particular value of VOIdx less than VOIdxMax) that belong to the target output views or that may be referred to through inter-view prediction in decoding depth view components belonging to the target output views. Such a picture parameter set RBSP is referred to as the active depth picture parameter set RBSP for the particular value of VOIdx. The restrictions on active picture parameter set RBSPs also apply to active depth picture parameter set RBSPs for a particular value of VOIdx less than VOIdxMax.

When a picture parameter set RBSP (with a particular value of pic_parameter_set_id) is not the active picture parameter set RBSP and it is referred to by a coded slice NAL unit belonging to a depth view component (i.e., with nal_unit_type equal to 21) and with VOIdx equal to VOIdxMax (using that value of pic_parameter_set_id), it is activated. This picture parameter set RBSP is called the active picture parameter set RBSP until it is deactivated when another picture parameter set RBSP becomes the active picture parameter set RBSP. A picture parameter set RBSP, with that particular value of pic_parameter_set_id, shall be available to the decoding process prior to its activation.

When a picture parameter set RBSP (with a particular value of `pic_parameter_set_id`) is not the active depth picture parameter set for a particular value of `VOIdx` less than `VOIdxMax` and it is referred to by a coded slice NAL unit belonging to a depth view component (i.e., with `nal_unit_type` equal to 21) and with the particular value of `VOIdx` (using that value of `pic_parameter_set_id`), it is activated for view components with the particular value of `VOIdx`. This picture parameter set RBSP is called the active depth picture parameter set RBSP for the particular value of `VOIdx` until it is deactivated when another picture parameter set RBSP becomes the active depth picture parameter set RBSP for the particular value of `VOIdx`. A picture parameter set RBSP, with that particular value of `pic_parameter_set_id`, shall be available to the decoding process prior to its activation.

When a picture parameter set RBSP (with a particular value of `pic_parameter_set_id`) is not the active texture picture parameter set for a particular value of `VOIdx` less than or equal to `VOIdxMax` and it is referred to by a coded slice NAL unit belonging to a texture view component (i.e., with `nal_unit_type` equal to 1, 5 or 20) and with the particular value of `VOIdx` (using that value of `pic_parameter_set_id`), it is activated for texture view components with the particular value of `VOIdx`. This picture parameter set RBSP is called the active texture picture parameter set RBSP for the particular value of `VOIdx` until it is deactivated when another picture parameter set RBSP becomes the active texture picture parameter set RBSP for the particular value of `VOIdx`. A picture parameter set RBSP, with that particular value of `pic_parameter_set_id`, shall be available to the decoding process prior to its activation.

Any picture parameter set NAL unit containing the value of `pic_parameter_set_id` for the active picture parameter set RBSP for a coded picture shall have the same content as that of the active picture parameter set RBSP for this coded picture unless it follows the last VCL NAL unit of this coded picture and precedes the first VCL NAL unit of another coded picture. Any picture parameter set NAL unit containing the value of `pic_parameter_set_id` for the active depth picture parameter set RBSP for a particular value of `VOIdx` less than `VOIdxMax` for a coded picture shall have the same content as that of the active view picture parameter set RBSP for the particular value of `VOIdx` for this coded picture unless it follows the last VCL NAL unit of this coded picture and precedes the first VCL NAL unit of another coded picture. Any picture parameter set NAL unit containing the value of `pic_parameter_set_id` for the active texture picture parameter set RBSP for a particular value of `VOIdx` for a coded picture shall have the same content as that of the active texture picture parameter set RBSP for the particular value of `VOIdx` for this coded picture unless it follows the last VCL NAL unit of this coded picture and precedes the first VCL NAL unit of another coded picture.

A MVCD sequence parameter set RBSP includes parameters that can be referred to by one or more picture parameter set RBSPs or one or more buffering period SEI messages.

Each MVCD sequence parameter set RBSP is initially considered not active at the start of the operation of the decoding process. At most one MVCD sequence parameter set RBSP is considered as the active MVCD sequence parameter set RBSP at any given moment during the operation of the decoding process, and when any particular MVCD sequence parameter set RBSP becomes the active MVCD sequence parameter set RBSP, the previously-active MVCD sequence parameter set RBSP (if any) is deactivated.

In addition to the active MVCD sequence parameter set RBSP, zero or more MVCD sequence parameter set RBSPs may be specifically active for view components (with a particular value of `VOIdx` less than `VOIdxMax`) that belong to the target output views or that may be referred to through inter-view prediction in decoding view components belonging to the target output views. Such a MVCD sequence parameter set RBSP is referred to as the active view MVCD sequence parameter set RBSP for the particular value of `VOIdx`. The restrictions on active MVCD sequence parameter set RBSPs also apply to active view MVCD sequence parameter set RBSPs for a particular value of `VOIdx` less than `VOIdxMax`.

Furthermore, zero or more MVCD sequence parameter set RBSPs may be specifically active for texture view components (with a particular value of `VOIdx` less than or equal to `VOIdxMax`) that belong to the target output views or that may be referred to through inter-view prediction in decoding texture view components belonging to the target output views. Such a MVCD sequence parameter set RBSP is referred to as the active texture MVCD sequence parameter set RBSP for the particular value of `VOIdx`. The restrictions on active MVCD sequence parameter set RBSPs also apply to active texture MVCD sequence parameter set RBSPs for a particular value of `VOIdx`.

For the following specification, the activating buffering period SEI message is specified as follows.

- If `VOIdxMax` is equal to `VOIdxMin` and the access unit contains a buffering period SEI message not included in an MVC scalable nesting SEI message and not included in a MVCD scalable nesting SEI message, this buffering period SEI message is the activating buffering period SEI message.

- Otherwise if `VOIDxMax` is not equal to `VOIDxMin` and the access unit contains a buffering period SEI message included in a MVCD scalable nesting SEI message and associated with the operation point being decoded, this buffering period SEI message is the activating buffering period SEI message.
- Otherwise, the access unit does not contain an activating buffering period SEI message.

When a sequence parameter set RBSP (`nal_unit_type` is equal to 7) with a particular value of `seq_parameter_set_id` is not already the active MVCD sequence parameter set RBSP and it is referred to by activation of a picture parameter set RBSP (using that value of `seq_parameter_set_id`) and the picture parameter set RBSP is activated by a coded slice NAL unit with `nal_unit_type` equal to 1 or 5 (the picture parameter set RBSP becomes the active picture parameter set RBSP and `VOIDxMax` is equal to `VOIDxMin` and there is no depth view component in any access unit) and the access unit does not contain an activating buffering period SEI message, it is activated. This sequence parameter set RBSP is called the active MVCD sequence parameter set RBSP until it is deactivated when another MVCD sequence parameter set RBSP becomes the active MVCD sequence parameter set RBSP. A sequence parameter set RBSP, with that particular value of `seq_parameter_set_id`, shall be available to the decoding process prior to its activation.

When a sequence parameter set RBSP (`nal_unit_type` is equal to 7) with a particular value of `seq_parameter_set_id` is not already the active MVCD sequence parameter set RBSP and it is referred to by an activating buffering period SEI message (using that value of `seq_parameter_set_id`) that is not included in a MVCD scalable nesting SEI message and `VOIDxMax` is equal to `VOIDxMin` and there is no depth view component in the access unit, it is activated. This sequence parameter set RBSP is called the active MVCD sequence parameter set RBSP until it is deactivated when another MVCD sequence parameter set RBSP becomes the active MVCD sequence parameter set RBSP. A sequence parameter set RBSP, with that particular value of `seq_parameter_set_id`, shall be available to the decoding process prior to its activation.

When a subset sequence parameter set RBSP (`nal_unit_type` is equal to 15) with a particular value of `seq_parameter_set_id` is not already the active MVCD sequence parameter set RBSP and it is referred to by activation of a picture parameter set RBSP (using that value of `seq_parameter_set_id`) and the picture parameter set RBSP is activated by a coded slice depth extension NAL unit with `nal_unit_type` equal to 21 and with `VOIDx` equal to `VOIDxMax` (the picture parameter set RBSP becomes the active picture parameter set RBSP) and the access unit does not contain an activating buffering period SEI message, it is activated. This subset sequence parameter set RBSP is called the active MVCD sequence parameter set RBSP until it is deactivated when another MVCD sequence parameter set RBSP becomes the active MVCD sequence parameter set RBSP. A subset sequence parameter set RBSP, with that particular value of `seq_parameter_set_id`, shall be available to the decoding process prior to its activation.

When a subset sequence parameter set RBSP (`nal_unit_type` is equal to 15) with a particular value of `seq_parameter_set_id` is not already the active MVCD sequence parameter set RBSP and it is referred to by an activating buffering period SEI message (using that value of `seq_parameter_set_id`) that is included in a MVCD scalable nesting SEI message, it is activated. This subset sequence parameter set RBSP is called the active MVCD sequence parameter set RBSP until it is deactivated when another MVCD sequence parameter set RBSP becomes the active MVCD sequence parameter set RBSP. A subset sequence parameter set RBSP, with that particular value of `seq_parameter_set_id`, shall be available to the decoding process prior to its activation.

NOTE 2 – The active MVCD sequence parameter set RBSP is either a sequence parameter set RBSP or a subset sequence parameter set RBSP. Sequence parameter set RBSPs are activated by coded slice NAL units with `nal_unit_type` equal to 1 or 5 or buffering period SEI messages that are not included in an MVC scalable nesting SEI message or a MVCD scalable nesting SEI message. Subset sequence parameter sets are activated by coded slice depth extension NAL units (`nal_unit_type` equal to 21) or buffering period SEI messages that are included in a MVCD scalable nesting SEI message. A sequence parameter set RBSP and a subset sequence parameter set RBSP may have the same value of `seq_parameter_set_id`.

For the following specification, the activating texture buffering period SEI message for a particular value of `VOIDx` is specified as follows.

- If the access unit contains one or more than one buffering period SEI message included in an MVC scalable nesting SEI message and associated with an operation point for which the greatest `VOIDx` in the associated bitstream subset is equal to the particular value of `VOIDx`, the first of these buffering period SEI messages, in decoding order, is the activating texture buffering period SEI message for the particular value of `VOIDx`.

- Otherwise, if the access unit contains a buffering period SEI message not included in an MVC scalable nesting SEI message or a MVCD scalable nesting SEI message, this buffering period SEI message is the activating texture buffering period SEI message for the particular value of VOIdx equal to VOIdxMin.
- Otherwise, the access unit does not contain an activating texture buffering period SEI message for the particular value of VOIdx.

When a sequence parameter set RBSP (nal_unit_type is equal to 7) with a particular value of seq_parameter_set_id is not already the active texture MVCD sequence parameter set RBSP for VOIdx equal to VOIdxMin and it is referred to by activation of a picture parameter set RBSP (using that value of seq_parameter_set_id) and the picture parameter set RBSP is activated by a coded slice NAL unit with nal_unit_type equal to 1 or 5 (the picture parameter set RBSP becomes the active texture picture parameter set RBSP for VOIdx equal to VOIdxMin), it is activated for texture view components with VOIdx equal to VOIdxMin. This sequence parameter set RBSP is called the active texture MVCD sequence parameter set RBSP for VOIdx equal to VOIdxMin until it is deactivated when another MVCD sequence parameter set RBSP becomes the active texture MVCD sequence parameter set RBSP for VOIdx equal to VOIdxMin. A sequence parameter set RBSP, with that particular value of seq_parameter_set_id, shall be available to the decoding process prior to its activation.

When a sequence parameter set RBSP (nal_unit_type is equal to 7) with a particular value of seq_parameter_set_id is not already the active texture MVCD sequence parameter set RBSP for VOIdx equal to VOIdxMin and it is referred to by an activating texture buffering period SEI message (using that value of seq_parameter_set_id) that is not included in an MVC scalable nesting SEI message or a MVCD scalable nesting SEI message, the sequence parameter set RBSP is activated for texture view components with VOIdx equal to VOIdxMin. This sequence parameter set RBSP is called the active texture MVCD sequence parameter set RBSP for VOIdx equal to VOIdxMin until it is deactivated when another MVCD sequence parameter set RBSP becomes the active texture MVCD sequence parameter set RBSP for VOIdx equal to VOIdxMin. A sequence parameter set RBSP, with that particular value of seq_parameter_set_id, shall be available to the decoding process prior to its activation.

When a subset sequence parameter set RBSP (nal_unit_type is equal to 15) with a particular value of seq_parameter_set_id is not already the active texture MVCD sequence parameter set RBSP for a particular value of VOIdx less than or equal to VOIdxMax and it is referred to by activation of a picture parameter set RBSP (using that value of seq_parameter_set_id) and the picture parameter set RBSP is activated by a coded slice MVC extension NAL unit (nal_unit_type equal to 20) with the particular value of VOIdx (the picture parameter set RBSP becomes the active texture picture parameter set RBSP for the particular value of VOIdx), it is activated for texture view components with the particular value of VOIdx. This subset sequence parameter set RBSP is called the active texture MVCD sequence parameter set RBSP for the particular value of VOIdx until it is deactivated when another MVCD sequence parameter set RBSP becomes the active texture MVCD sequence parameter set RBSP for the particular value of VOIdx. A subset sequence parameter set RBSP, with that particular value of seq_parameter_set_id, shall be available to the decoding process prior to its activation.

When a subset sequence parameter set RBSP (nal_unit_type is equal to 15) with a particular value of seq_parameter_set_id is not already the active texture MVCD sequence parameter set RBSP for a particular value of VOIdx less than or equal to VOIdxMax and it is referred to by an activating texture buffering period SEI message (using that value of seq_parameter_set_id) that is included in an MVC scalable nesting SEI message and associated with the particular value of VOIdx, this subset sequence parameter set RBSP is activated for texture view components with the particular value of VOIdx. This subset sequence parameter set RBSP is called the active texture MVCD sequence parameter set RBSP for the particular value of VOIdx until it is deactivated when another MVCD sequence parameter set RBSP becomes the active texture MVCD sequence parameter set RBSP for the particular value of VOIdx. A subset sequence parameter set RBSP, with that particular value of seq_parameter_set_id, shall be available to the decoding process prior to its activation.

For the following specification, the activating view buffering period SEI message for a particular value of VOIdx is specified as follows.

- If the access unit contains one or more than one buffering period SEI message included in a MVCD scalable nesting SEI message and associated with an operation point for which the greatest VOIdx in the associated bitstream subset is equal to the particular value of VOIdx, the first of these buffering period SEI messages, in decoding order, is the activating view buffering period SEI message for the particular value of VOIdx.

- Otherwise, the access unit does not contain an activating view buffering period SEI message for the particular value of `VOIDx`.

When a subset sequence parameter set RBSP (`nal_unit_type` is equal to 15) with a particular value of `seq_parameter_set_id` is not already the active view MVCD sequence parameter set RBSP for a particular value of `VOIDx` less than `VOIDxMax` and it is referred to by activation of a picture parameter set RBSP (using that value of `seq_parameter_set_id`) and the picture parameter set RBSP is activated by a coded slice NAL unit with `nal_unit_type` equal to 21 and with the particular value of `VOIDx` (the picture parameter set RBSP becomes the active view picture parameter set RBSP for the particular value of `VOIDx`), it is activated for view components with the particular value of `VOIDx`. This subset sequence parameter set RBSP is called the active view MVCD sequence parameter set RBSP for the particular value of `VOIDx` until it is deactivated when another MVCD sequence parameter set RBSP becomes the active view MVCD sequence parameter set RBSP for the particular value of `VOIDx` or when decoding an access unit with `VOIDxMax` less than or equal to the particular value of `VOIDx`. A subset sequence parameter set RBSP, with that particular value of `seq_parameter_set_id`, shall be available to the decoding process prior to its activation.

When a subset sequence parameter set RBSP (`nal_unit_type` is equal to 15) with a particular value of `seq_parameter_set_id` is not already the active view MVCD sequence parameter set RBSP for a particular value of `VOIDx` less than `VOIDxMax` and it is referred to by an activating view buffering period SEI message (using that value of `seq_parameter_set_id`) that is included in a MVCD scalable nesting SEI message and associated with the particular value of `VOIDx`, this subset sequence parameter set RBSP is activated for view components with the particular value of `VOIDx`. This subset sequence parameter set RBSP is called the active view MVCD sequence parameter set RBSP for the particular value of `VOIDx` until it is deactivated when another MVCD sequence parameter set RBSP becomes the active view MVCD sequence parameter set RBSP for the particular value of `VOIDx` or when decoding an access unit with `VOIDxMax` less than or equal to the particular value of `VOIDx`. A subset sequence parameter set RBSP, with that particular value of `seq_parameter_set_id`, shall be available to the decoding process prior to its activation.

A MVCD sequence parameter set RBSP that includes a value of `profile_idc` not specified in Annex A or Annex H or Annex I shall not be referred to by activation of a picture parameter set RBSP as the active picture parameter set RBSP or as active view picture parameter set RBSP or as active texture picture parameter set RBSP (using that value of `seq_parameter_set_id`) or referred to by a buffering period SEI message (using that value of `seq_parameter_set_id`). A MVCD sequence parameter set RBSP including a value of `profile_idc` not specified in Annex A or Annex H or Annex I is ignored in the decoding for profiles specified in Annex A or Annex H or Annex I.

It is a requirement of bitstream conformance that the following constraints are obeyed:

- For each particular value of `VOIDx`, all coded slice NAL units (with `nal_unit_type` equal to 1, 5, 20, or 21) of a coded video sequence shall refer to the same value of `seq_parameter_set_id` (via the picture parameter set RBSP that is referred to by the value of `pic_parameter_set_id`).
- The value of `seq_parameter_set_id` in a buffering period SEI message that is not included in an MVC scalable nesting SEI message shall be identical to the value of `seq_parameter_set_id` in the picture parameter set RBSP that is referred to by coded slice NAL units with `nal_unit_type` equal to 1 or 5 (via the value of `pic_parameter_set_id`) in the same access unit.
- The value of `seq_parameter_set_id` in a buffering period SEI message that is included in an MVC scalable nesting SEI message and is associated with a particular value of `VOIDx` shall be identical to the value of `seq_parameter_set_id` in the picture parameter set RBSP that is referred to by coded slice NAL units with `nal_unit_type` equal to 1, 5 or 20 with the particular value of `VOIDx` (via the value of `pic_parameter_set_id`) in the same access unit.
- The value of `seq_parameter_set_id` in a buffering period SEI message that is included in a MVCD scalable nesting SEI message and is associated with a particular value of `VOIDx` shall be identical to the value of `seq_parameter_set_id` in the picture parameter set RBSP that is referred to by coded slice NAL units with `nal_unit_type` equal to 21 with the particular value of `VOIDx` (via the value of `pic_parameter_set_id`) in the same access unit.

The active view MVCD sequence parameter set RBSPs for different values of `VOIDx` may be the same MVCD sequence parameter set RBSP. The active MVCD sequence parameter set RBSP and an active view MVCD sequence parameter set RBSP for a particular value of `VOIDx` may be the same MVCD sequence parameter set RBSP.

The active texture MVCD sequence parameter set RBSPs for different values of VOIdx may be the same MVCD sequence parameter set RBSP. The active MVCD sequence parameter set RBSP and an active texture MVCD sequence parameter set RBSP for a particular value of VOIdx may be the same MVCD sequence parameter set RBSP..

When the active MVCD sequence parameter set RBSP for a coded picture is a sequence parameter set RBSP, any sequence parameter set RBSP in the coded video sequence containing this coded picture and with the value of seq_parameter_set_id for the active MVCD sequence parameter set RBSP shall have the same content as that of the active MVCD sequence parameter set RBSP.

When the active MVCD sequence parameter set RBSP for a coded picture is a subset sequence parameter set RBSP, any subset sequence parameter set RBSP in the coded video sequence containing this coded picture and with the value of seq_parameter_set_id for the active MVCD sequence parameter set RBSP shall have the same content as that of the active MVCD sequence parameter set RBSP.

For each particular value of VOIdx, the following applies:

- When the active texture MVCD sequence parameter set RBSP for a coded picture is a sequence parameter set RBSP, any sequence parameter set RBSP in the coded video sequence containing this coded picture and with the value of seq_parameter_set_id for the active texture MVCD sequence parameter set RBSP shall have the same content as that of the active texture MVCD sequence parameter set RBSP.
- When the active texture MVCD sequence parameter set RBSP for a coded picture is a subset sequence parameter set RBSP, any subset sequence parameter set RBSP in the coded video sequence containing this coded picture and with the value of seq_parameter_set_id for the active texture MVCD sequence parameter set RBSP shall have the same content as that of the active texture MVCD sequence parameter set RBSP.
- The active view MVCD sequence parameter set RBSP for a coded picture is a subset sequence parameter set RBSP, and any subset sequence parameter set RBSP in the coded video sequence containing this coded picture and with the value of seq_parameter_set_id for the active view MVCD sequence parameter set RBSP shall have the same content as that of the active view MVCD sequence parameter set RBSP.

NOTE 3 – If picture parameter set RBSPs or MVCD sequence parameter set RBSPs are conveyed within the bitstream, these constraints impose an order constraint on the NAL units that contain the picture parameter set RBSPs or MVCD sequence parameter set RBSPs, respectively. Otherwise (picture parameter set RBSPs or MVCD sequence parameter set RBSPs are conveyed by other means not specified in this Recommendation | International Standard), they must be available to the decoding process in a timely fashion such that these constraints are obeyed.

When present, a sequence parameter set extension RBSP includes parameters having a similar function to those of a sequence parameter set RBSP. For purposes of establishing constraints on the syntax elements of the sequence parameter set extension RBSP and for purposes of determining activation of a sequence parameter set extension RBSP, the sequence parameter set extension RBSP shall be considered part of the preceding sequence parameter set RBSP with the same value of seq_parameter_set_id. When a sequence parameter set RBSP is present that is not followed by a sequence parameter set extension RBSP with the same value of seq_parameter_set_id prior to the activation of the sequence parameter set RBSP, the sequence parameter set extension RBSP and its syntax elements shall be considered not present for the active MVCD sequence parameter set RBSP. The contents of sequence parameter set extension RBSPs only apply when the base texture view, which conforms to one or more of the profiles specified in Annex A, of a coded video sequence conforming to one or more profiles specified in Annex I is decoded. Subset sequence parameter set RBSPs shall not be followed by a sequence parameter set extension RBSP.

NOTE 4 – Sequence parameter sets extension RBSPs are not considered to be part of a subset sequence parameter set RBSP and subset sequence parameter set RBSPs must not be followed by a sequence parameter set extension RBSP.

For view components with VOIdx equal to VOIdxMax, all constraints that are expressed on the relationship between the values of the syntax elements (and the values of variables derived from those syntax elements) in MVCD sequence parameter sets and picture parameter sets and other syntax elements are expressions of constraints that apply only to the active MVCD sequence parameter set and the active picture parameter set. For view components with a particular value of VOIdx less than VOIdxMax, all constraints that are expressed on the relationship between the values of the syntax elements (and the values of variables derived from those syntax elements) in MVCD sequence parameter sets and picture

parameter sets and other syntax elements are expressions of constraints that apply only to the active view MVCD sequence parameter set and the active view picture parameter set for the particular value of `VOIDx`. If any MVCD sequence parameter set RBSP having `profile_idc` equal to the value of one of the `profile_idc` values specified in Annex A or Annex H or Annex I is present that is never activated in the bitstream (i.e., it never becomes the active MVCD sequence parameter set or an active view MVCD sequence parameter set), its syntax elements shall have values that would conform to the specified constraints if it were activated by reference in an otherwise-conforming bitstream. If any picture parameter set RBSP is present that is never activated in the bitstream (i.e., it never becomes the active picture parameter set or an active view picture parameter set), its syntax elements shall have values that would conform to the specified constraints if it were activated by reference in an otherwise-conforming bitstream.

During operation of the decoding process (see subclause I.8), for view components with `VOIDx` equal to `VOIDxMax`, the values of parameters of the active picture parameter set and the active MVCD sequence parameter set shall be considered in effect. For view components with a particular value of `VOIDx` less than `VOIDxMax`, the values of the parameters of the active view picture parameter set and the active view MVCD sequence parameter set for the particular value of `VOIDx` shall be considered in effect. For interpretation of SEI messages that apply to the entire access unit or the view component with `VOIDx` equal to `VOIDxMax`, the values of the parameters of the active picture parameter set and the active MVCD sequence parameter set for the same access unit shall be considered in effect unless otherwise specified in the SEI message semantics. For interpretation of SEI messages that apply to view components with a particular value of `VOIDx` less than `VOIDxMax`, the values of the parameters of the active view picture parameter set and the active view MVCD sequence parameter set for the particular value of `VOIDx` for the same access unit shall be considered in effect unless otherwise specified in the SEI message semantics.

For any active MVCD sequence parameter set or active view MVCD sequence parameter set, part of the syntax elements in the MVC sequence parameter set extension applies only to the depth views referring to this sequence parameter set, while the some other parts of the syntax elements in the MVCD sequence parameter set extension collectively apply to both the depth views referring to this sequence parameter set and the corresponding texture views. More specifically, the view dependency information of the MVCD sequence parameter set extension applies only to the depth views, and the level definitions collectively apply to operation points, each of which contains both depth views and their corresponding texture views. Moreover, the `mvcd_vui_parameters_extension()` applies collectively to both the depth views referring to this MVCD sequence parameter set and the corresponding texture views. The `vui_parameters()` included in the sequence parameter set data syntax structure, if present, apply collectively to both the depth views referring to this sequence parameter set and the corresponding texture views, except for the aspect ratio information and the bitstream restriction information, if present, which apply only to the depth views referring to this MVCD sequence parameter set. The aspect ratio information and the bitstream restriction information for the texture views may be present in the `vui_parameters()` syntax structure included in an MVC sequence parameter set.

I.7.4.1.2.2 Order of access units and association to coded video sequences

The specification of subclause H.7.4.1.2.2 apply.

I.7.4.1.2.3 Order of NAL units and coded pictures and association to access units

The specification of subclause H.7.4.1.2.3 applies with the following modifications.

NOTE – Some bitstreams that conform to one or more profiles specified in this annex do not conform to any profile specified in Annex A (prior to operation of the base view extraction process specified in subclause I.8.5.4). As specified in subclauses 7.4.1 and 7.4.1.2.3 for the profiles specified in Annex A, NAL units with `nal_unit_type` equal to 21 are classified as non-VCL NAL units that must be preceded within each access unit by at least one NAL unit with `nal_unit_type` in the range of 1 to 5, inclusive. For this reason, any bitstream that conforms to one or more profiles specified in this annex does not conform to any profile specified in Annex A when it contains any of the following:

- Any access unit that does not contain any NAL units with `nal_unit_type` equal to 1 or 5, but contains one or more NAL units with `nal_unit_type` equal to 6, 7, 8, 9, or 15.
- Any access unit in which one or more NAL units with `nal_unit_type` equal to 7, 8, or 15 is present after the last NAL unit in the access unit with `nal_unit_type` equal to 1 or 5.

The association of VCL NAL units to primary or redundant coded pictures is specified in subclause I.7.4.1.2.5.

The constraints for the detection of the first VCL NAL unit of a primary coded picture are specified in subclause I.7.4.1.2.4.

The constraint expressed in subclause H.7.4.1.2.3 on the order of a buffering period SEI message is replaced by the following constraints.

- When an SEI NAL unit containing a buffering period SEI message is present, the following applies:
 - If the buffering period SEI message is the only buffering period SEI message in the access unit and it is not included in an MVC scalable nesting SEI message or a MVCD scalable nesting SEI message, the buffering period SEI message shall be the first SEI message payload of the first SEI NAL unit in the access unit.
 - Otherwise (the buffering period SEI message is not the only buffering period SEI message in the access unit or it is included in an MVC scalable nesting SEI message or it is included in a MVCD scalable nesting SEI message), the following constraints are specified:
 - When a buffering period SEI message that is not included in either an MVC scalable nesting SEI message or a MVCD scalable nesting SEI message is present, this buffering period SEI message shall be the only SEI message payload of the first SEI NAL unit in the access unit.
 - An MVC scalable nesting SEI message that includes a buffering period SEI message shall not include any other SEI messages and shall be the only SEI message inside the SEI NAL unit.
 - A MVCD scalable nesting SEI message that includes a buffering period SEI message shall not include any other SEI messages and shall be the only SEI message inside the SEI NAL unit.
 - All SEI NAL units that precede an SEI NAL unit that contains an MVC scalable nesting SEI message with a buffering period SEI message as payload, or a MVCD scalable nesting SEI message with a buffering period SEI message as payload in an access unit shall only contain buffering period SEI messages or MVC scalable nesting SEI messages with a buffering period SEI message as payload, or MVCD scalable nesting SEI messages with a buffering period SEI message.

I.7.4.1.2.4 Detection of the first VCL NAL unit of a primary coded picture

The specification of subclause H.7.4.1.2.4 applies.

I.7.4.1.2.5 Order of VCL NAL units and association to coded pictures

The specification of subclause H.7.4.1.2.5 applies with following modifications.

Each VCL NAL unit is part of a coded picture.

Let voidx be the value of VOIDx of any particular VCL NAL unit. The order of the VCL NAL units within a coded picture is constrained as follows:

- For all VCL NAL units following this particular VCL NAL unit, the value of VOIDx shall be greater than or equal to voidx .
- All VCL NAL units for a depth view component, if present, shall follow any VCL NAL unit of a texture view component with a same value of VOIDx .

For each set of VCL NAL units within a texture or depth view component, the following applies:

- If arbitrary slice order, as specified in Annex A, subclause H.10 or subclause I.10, is allowed, coded slice NAL units of a view component may have any order relative to each other.
- Otherwise (arbitrary slice order is not allowed), coded slice NAL units of a slice group shall not be interleaved with coded slice NAL units of another slice group and the order of coded slice NAL units within a slice group shall be in the order of increasing macroblock address for the first macroblock of each coded slice NAL unit of the same slice group.

The following applies:

- If a coded texture view component with a particular view_id is the first field view component of a complementary field pair, the depth view component with the same view_id value, if present in the access unit, shall be a coded frame view component or the first field view component of a complementary field pair.
- Otherwise, if a coded texture view component with a particular view_id is the second field view component of a complementary field pair, the depth view component with the same view_id value, if present in the access unit, shall be the second field view component of a complementary field pair.
- Otherwise, if a coded texture view component with a particular view_id is a non-paired field, the depth view component with the same view_id value, if present in the access unit, shall be a coded frame view component or a non-paired field.
- Otherwise (a coded texture view component with a particular view_id is a coded frame), the depth view component with the same view_id value, if present in the access unit, shall be a coded frame view component.

NAL units having nal_unit_type equal to 12 may be present in the access unit but shall not precede the first VCL NAL unit of the primary coded picture within the access unit.

NAL units having nal_unit_type equal to 0 or in the range of 24 to 31, inclusive, which are unspecified, may be present in the access unit but shall not precede the first VCL NAL unit of the primary coded picture within the access unit.

NAL units having nal_unit_type in the range of 22 to 23, inclusive, which are reserved, shall not precede the first VCL NAL unit of the primary coded picture within the access unit (when specified in the future by ITU-T | ISO/IEC).

I.7.4.2 Raw byte sequence payloads and RBSP trailing bits semantics

I.7.4.2.1 Sequence parameter set RBSP semantics

The semantics specified in subclause 7.4.2.1 apply.

I.7.4.2.1.1 Sequence parameter set data semantics

The semantics specified in subclause H.7.4.2.1.1 apply with the substitution of MVCD sequence parameter set for MVC sequence parameter set. All constraints specified in subclause H.7.4.2.1.1 apply only to the texture view components for which the MVCD sequence parameter set is the active texture MVC sequence parameter set or to the depth view components for which the MVCD sequence parameter set is the active view MVC sequence parameter set as specified in subclause I.7.4.1.2.1.

I.7.4.2.1.1.1 Scaling list semantics

The semantics specified in subclause H.7.4.2.1.1.1 apply.

I.7.4.2.1.2 Sequence parameter set extension RBSP semantics

The semantics specified in subclause 7.4.2.1.2 apply. Additionally, the following applies.

Sequence parameter set extension RBSPs can only follow sequence parameter set RBSPs in decoding order. Subset sequence parameter set RBSPs shall not be followed by a sequence parameter set extension RBSP. The contents of sequence parameter set extension RBSPs only apply when the base view, which conforms to one or more of the profiles specified in Annex A, of a coded video sequence conforming to one or more profiles specified in Annex I is decoded.

I.7.4.2.1.3 Subset sequence parameter set RBSP semantics

The semantics specified in subclause 7.4.2.1.3 apply with the following additions.

mvcd_vui_parameters_present_flag equal to 0 specifies that the syntax structure `mvcd_vui_parameters_extension()` corresponding to MVCD VUI parameters extension is not present. **mvcd_vui_parameters_present_flag** equal to 1 specifies that the syntax structure `mvcd_vui_parameters_extension()` is present and referred to as MVCD VUI parameters extension.

texture_vui_parameters_present_flag equal to 0 specifies that the syntax structure `mvcd_vui_parameters_extension()` corresponding to MVCD texture sub-bitstream VUI parameters extension is not present. **texture_vui_parameters_present_flag** equal to 1 specifies that the syntax structure `mvcd_vui_parameters_extension()` is present and referred to as MVCD texture sub-bitstream VUI parameters extension.

I.7.4.2.1.4 Sequence parameter set MVCD extension semantics

The semantics specified in subclause H.7.4.2.1.4 apply with the substitution of texture view component or depth view component for view component and with the following additions:

depth_view_present_flag[i] equal to 0 specifies that there is no depth view having a `view_id` equal to `view_id[i]` and `VOIdx` equal to `i`. **depth_view_present_flag[i]** equal to 1 specifies that there is a depth view having a `view_id` equal to `view_id[i]`.

texture_view_present_flag[i] equal to 0 specifies that there is no texture view having a `view_id` equal to `view_id[i]` and `VOIdx` equal to `i`. **texture_view_present_flag[i]** equal to 1 specifies that there is a texture view having a `view_id` equal to `view_id[i]` and `VOIdx` equal to `i`. When **depth_view_present_flag[i]** is equal to 0, **texture_view_present_flag[i]** shall be equal to 1.

`num_anchor_refs_l0[i]`, `anchor_ref_l0[i][j]`, `num_anchor_refs_l1[i]`, `anchor_ref_l1[i][j]`, `num_non_anchor_refs_l0[i]`, `non_anchor_ref_l0[i][j]`, `num_non_anchor_refs_l1[i]`, and `non_anchor_ref_l1[i][j]` apply to depth view components.

applicable_op_depth_flag[i][j][k] equal to 0 indicates that the depth view with `view_id` equal to `applicable_op_target_view_id[i][j][k]` is not included in the `j`-th operation point. **applicable_op_depth_flag[i][j][k]** equal to 1 indicates that the depth view with `view_id` equal to `applicable_op_target_view_id[i][j][k]` is included in the `j`-th operation point.

applicable_op_texture_flag[i][j][k] equal to 0 indicates that the texture view with `view_id` equal to `applicable_op_target_view_id[i][j][k]` is not included in the `j`-th operation point. **applicable_op_texture_flag[i][j][k]** equal to 1 indicates that the texture view with `view_id` equal to `applicable_op_target_view_id[i][j][k]` is included in the `j`-th operation point. When **applicable_op_depth_flag[i][j][k]** is equal to 0, **applicable_op_texture_flag[i][j][k]** shall be equal to 1.

applicable_op_num_texture_views_minus1[i][j] plus 1 specifies the number of texture views required for decoding the target output views corresponding to the `j`-th operation point to which the level indicated by `level_idc[i]` applies. The number of texture views specified by **applicable_op_num_views_minus1** includes the texture views of the target output views and the texture views that the target output views depend on. The value of **applicable_op_num_texture_views_minus1[i][j]** shall be in the range of 0 to 1023, inclusive.

applicable_op_num_depth_views[i][j] specifies the number of depth views required for decoding the target output views corresponding to the `j`-th operation point to which the level indicated by `level_idc[i]` applies. The number of depth views specified by **applicable_op_num_depth_views_minus1** includes the depth views of the target output views and the depth views that the depth views of the target output views depend on. The value of **applicable_op_num_depth_views_minus1[i][j]** shall be in the range of 0 to 1023, inclusive.

All sequence parameter set MVCD extensions that are included in the active view MVCD sequence parameter set RBSPs of one coded video sequence shall be identical.

I.7.4.2.2 Picture parameter set RBSP semantics

The semantics specified in subclause H.7.4.2.2 apply with substituting MVCD sequence parameter set for MVC sequence parameter set. All constraints specified in subclause H.7.4.2.2 apply only to the texture or depth view components for which the picture parameter set is the active picture parameter set or the active view picture parameter set

or the active texture picture parameter set as specified in subclause I.7.4.1.2.1.

I.7.4.2.3 Supplemental enhancement information RBSP semantics

The semantics specified in subclause H.7.4.2.3 apply.

I.7.4.2.3.1 Supplemental enhancement information message semantics

The semantics specified in subclause H.7.4.2.3.1 apply.

I.7.4.2.4 Access unit delimiter RBSP semantics

The semantics specified in subclause H.7.4.2.4 apply.

NOTE – The value of `primary_pic_type` applies to the `slice_type` values in all slice headers of the primary coded picture, including the `slice_type` syntax elements in all NAL units with `nal_unit_type` equal to 1, 5, 20 or 21. NAL units with `nal_unit_type` equal to 2 are not present in bitstreams conforming to any of the profiles specified in this annex.

I.7.4.2.5 End of sequence RBSP semantics

The semantics specified in subclause H.7.4.2.5 apply.

I.7.4.2.6 End of stream RBSP semantics

The semantics specified in subclause H.7.4.2.6 apply.

I.7.4.2.7 Filler data RBSP semantics

The semantics specified in subclause H.7.4.2.7 apply with the following modifications.

Filler data NAL units shall be considered to contain the syntax elements `priority_id`, `view_id`, and `temporal_id` with values that are inferred as follows:

- Let `prevMvcNalUnit` be the most recent NAL unit in decoding order that has `nal_unit_type` equal to 14, 20 or 21.
NOTE – The most recent NAL unit in decoding order with `nal_unit_type` equal to 14, 20 or 21 always belongs to the same access unit as the filler data NAL unit.
- The values of `priority_id`, `view_id`, and `temporal_id` for the filler data NAL unit are inferred to be equal to the values of `priority_id`, `view_id`, and `temporal_id`, respectively, of the NAL unit `prevMvcNalUnit`.

I.7.4.2.8 Slice layer without partitioning RBSP semantics

The semantics specified in subclause H.7.4.2.8 apply.

I.7.4.2.9 Slice data partition RBSP semantics

Slice data partition syntax is not present in bitstreams conforming to one or more of the profiles specified in Annex I.

I.7.4.2.10 RBSP slice trailing bits semantics

The semantics specified in subclause H.7.4.2.10 apply.

I.7.4.2.11 RBSP trailing bits semantics

The semantics specified in subclause H.7.4.2.11 apply.

I.7.4.2.12 Prefix NAL unit RBSP semantics

The semantics specified in subclause H.7.4.2.12 apply.

I.7.4.2.13 Slice layer extension RBSP semantics

The semantics specified in subclause H.7.4.2.13 apply.

I.7.4.3 Slice header semantics

The semantics specified in subclause H.7.4.3 apply with the substitution of texture view component (for `nal_unit_type` equal to 1, 5, and 20) or depth view component (for `nal_unit_type` equal to 21) for view component and with the following modifications.

When `nal_unit_type` is equal to 1, 5, or 20, all constraints specified in subclause H.7.4.3 apply only to the texture view components with the same value of `VOIDx`. When `nal_unit_type` is equal to 21, all constraints specified in subclause H.7.4.3 apply only to the depth view components with the same value of `VOIDx`.

The value of the following MVCD sequence parameter set syntax elements shall be the same across all coded slice NAL units of `nal_unit_type` equal to 1, 5, and 20 of an access unit: `chroma_format_idc`.

The value of the following slice header syntax elements shall be the same across all coded slice NAL units of `nal_unit_type` equal to 1, 5, and 20 of an access unit: `field_pic_flag` and `bottom_field_flag`.

The value of the following slice header syntax elements shall be the same across all coded slice NAL units of `nal_unit_type` equal to 21 of an access unit: `field_pic_flag` and `bottom_field_flag`.

I.7.4.3.1 Reference picture list modification semantics

The semantics specified in subclause H.7.4.3.1 apply.

I.7.4.3.1.1 Reference picture list MVC modification semantics

The semantics specified in subclause H.7.4.3.1 apply.

I.7.4.3.2 Prediction weight table semantics

The semantics specified in subclause H.7.4.3.2 apply.

I.7.4.3.3 Decoded reference picture marking semantics

The semantics specified in subclause 7.4.3.3 apply to each view independently, with "sequence parameter set" being replaced by "MVCD sequence parameter set", and "primary coded picture" being replaced by "texture view component" for `nal_unit_type` equal to 1, 5, and 20, and by "depth view component" for `nal_unit_type` equal to 21.

I.7.4.4 Slice data semantics

The semantics specified in subclause H.7.4.4 apply.

I.7.4.5 Macroblock layer semantics

The semantics specified in subclause H.7.4.5 apply.

I.7.4.5.1 Macroblock prediction semantics

The semantics specified in subclause H.7.4.5.1 apply.

I.7.4.5.2 Sub-macroblock prediction semantics

The semantics specified in subclause H.7.4.5.2 apply.

I.7.4.5.3 Residual data semantics

The semantics specified in subclause H.7.4.5.3 apply.

I.7.4.5.3.1 Residual luma semantics

The semantics specified in subclause H.7.4.5.3.1 apply.

I.7.4.5.3.2 Residual block CAVLC semantics

The semantics specified in subclause H.7.4.5.3.2 apply.

I.7.4.5.3.3 Residual block CABAC semantics

The semantics specified in subclause H.7.4.5.3.3 apply.

I.8 MVCD decoding process

This subclause specifies the decoding process for an access unit of a coded video sequence conforming to one or more of the profiles specified in Annex I. Specifically, this subclause specifies how the decoded picture with multiple texture view components and multiple depth view components is derived from syntax elements and global variables that are derived from NAL units in an access unit when the decoder is decoding the operation point identified by the target temporal level and the target output texture and depth views.

The decoding process is specified such that all decoders shall produce numerically identical results for the target output texture and depth views. Any decoding process that produces identical results for the target output texture and depth views to the process described here conforms to the decoding process requirements of this Recommendation | International Standard.

Unless stated otherwise, the syntax elements and derived upper-case variables that are referred to by the decoding process specified in this subclause and all child processes invoked from the process specified in this subclause are the syntax elements and derived upper-case variables for the current access unit.

The target output texture and depth views are either specified by external means not specified in this Specification, or, when not specified by external means, there shall be one target output texture view which is the base texture view.

NOTE – The association of VOIdx values to view_id values according to the decoding process of clause I.8 may differ from that of the decoding process of clause H.8.

A target output view may include only a texture view, only a depth view, or both the texture view and the depth view, which have the same view_id value.

All sub-bitstreams that can be derived using the sub-bitstream extraction process with depthPresentFlagTarget equal to 0 or 1, pIdTarget equal to any value in the range of 0 to 63, inclusive, tIdTarget equal to any value in the range of 0 to 7, inclusive, viewIdTargetList consisting of any one or more viewIdTarget's identifying the views in the bitstream as inputs as specified in subclause I.8.5 shall result in a set of coded video sequences, with each coded video sequence conforming to one or more of the profiles specified in Annex A, Annex H and Annex I.

Let vOIdxList be a list of integer values specifying the VOIdx values of the view components of the access unit. The variable VOIdxMax is set equal to the maximum value of the entries in the list vOIdxList, and the variable vOIdxMin is set to the minimum value of the entries in the list vOIdxList. When the current access unit is an anchor access unit, the variable VOIdxMin is set to vOIdxMin.

The MVCD video decoding process specified in this subclause is repeatedly invoked for each texture and depth view component with VOIdx from vOIdxMin to VOIdxMax, inclusive, which is present in the list vOIdxList, in increasing order of VOIdx and in decoding order of texture or depth view components as specified in subclause I.7.4.1.2.5.

Outputs of the MVCD video decoding process are decoded samples of the current primary coded picture including all decoded texture and depth view components of the target output texture and depth views.

For each texture view component and each depth view component, the specifications in clause H.8 apply, with the decoding processes for picture order count, reference picture lists construction and decoded reference picture marking being modified in subclauses I.8.1, I.8.2, I.8.3, and I.8.4, respectively. The MVCD inter prediction and inter-view prediction process is specified in subclause I.8.4.

I.8.1 MVCD decoding process for picture order count

The specifications in subclause 8.2.1 apply independently for each texture view or depth view.

I.8.2 MVC decoding process for reference picture lists construction

The specification of subclause H.8.2 apply with substituting "view component" as either "texture view component" or "depth view component", and "frame view component" as either "depth frame view component" or "texture frame view component", and "field view component" as "texture field view component" or "depth field view component".

Additionally, an inter-view reference component or the inter-view only reference component is identified by the view_id and a depth view component when the current slice is a part of a coded depth view component or a texture view component if the current slice is a part of a coded texture view component.

I.8.2.1 Initialisation process for reference picture list for inter-view prediction references

The specifications of subclause H.8.2.1 apply.

I.8.2.2 Modification process for reference picture lists

The specifications of subclause H.8.2.2 apply.

I.8.2.2.1 Modification process of reference picture lists for short-term reference pictures for inter prediction

The specifications of subclause H.8.2.2.1 apply.

I.8.2.2.2 Modification process of reference picture lists for long-term reference pictures for inter prediction

The specifications of subclause H.8.2.2.2 apply.

I.8.2.2.3 Modification process for reference picture lists for inter-view prediction references

The specifications of subclause H.8.2.2.3 apply.

I.8.3 MVCD decoded reference picture marking process

The specifications of subclause H.8.3 apply. Additionally, following applies.

The process specified in this subclause is invoked for a particular texture view or depth view with view order index VOIdx. The specifications in subclause H.8.3 apply with "view component" being replaced by either "texture view component" or "depth view component", "frame view component" being replaced by either "texture frame view component" or "depth frame view component", and "field view component" being replaced by either "texture field view component" or "depth field view component". During the invocation of the process for a particular texture view, only texture view components of the particular view are considered. During the invocation of the process for a particular depth

view, only depth view components of the particular view are considered. The marking of view components of other views is not changed.

NOTE – A texture view component of a picture may have a different marking status than other texture view components of the same picture. A depth view component of a picture may have a different marking status than other depth view components of the same picture. A texture view component of a picture may have a different marking status than a depth view component.

I.8.4 MVCD inter prediction and inter-view prediction process

The specifications of subclause H.8.4 apply.

I.8.5 Specification of bitstream subsets

The specifications of subclause H.8.5 apply.

I.8.5.1 Derivation process for required anchor view components

The specification of subclause H.8.5.1 apply with substituting "view component" with "depth view component" and "view" with "depth view" or and substituting "required for anchor" with "required for anchor depth".

The specification of subclause H.8.5.1 apply with substituting "view component" with "texture view component" and "view" with "texture view" and substituting "required for anchor" with "required for anchor texture".

I.8.5.2 Derivation process for required non-anchor view components

The specification of subclause H.8.5.2 apply with substituting "view component" with "depth view component" and "view" with "depth view" and substituting "required for anchor" with "required for anchor depth".

The specification of subclause H.8.5.2 apply with substituting "view component" with "depth view component" with either "depth view", and substituting "required for anchor" with "required for anchor texture".

I.8.5.3 Sub-bitstream extraction process

It is requirement of bitstream conformance that any sub-bitstream that is the output of the process specified in this subclause with depthPresentFlagTarget equal to 0 or 1, pIdTarget equal to any value in the range of 0 to 63, inclusive, tIdTarget equal to any value in the range of 0 to 7, inclusive, viewIdTargetList consisting of any one or more values of viewIdTarget identifying the views in the bitstream, shall be conforming to this Recommendation | International Standard.

NOTE 1 – A conforming bitstream contains one or more coded slice NAL units with priority_id equal to 0 and temporal_id equal to 0.

NOTE 2 – It is possible that not all operation points of sub-bitstreams resulting from the sub-bitstream extraction process have an applicable level_id or level_idc[i]. In this case, each coded video sequence in a sub-bitstream must still conform to one or more of the profiles specified in Annex A, Annex H and Annex I, but may not satisfy the level constraints specified in subclauses A.3, H.10.2 and I.10.2, respectively.

Inputs to this process are:

- a variable depthPresentFlagTarget (when present),
- a variable pIdTarget (when present),
- a variable tIdTarget (when present),
- a list viewIdTargetList consisting of one or more values of viewIdTarget (when present).
- a list viewIdDepthTargetList consisting of one or more value of viewIdDepthTarget (when present).

Outputs of this process are a sub-bitstream and a list of VOIdx values VOIdxList.

When depthPresentFlagTarget is not present as input to this subclause, depthPresentFlagTarget is inferred to be equal to 0.

When pIdTarget is not present as input to this subclause, pIdTarget is inferred to be equal to 63.

When tIdTarget is not present as input to this subclause, tIdTarget is inferred to be equal to 7.

When viewIdTargetList is not present as input to this subclause, there shall be one value of viewIdTarget inferred in viewIdTargetList and the value of viewIdTarget is inferred to be equal to view_id of the base view.

When viewIdDepthTargetList is not present as input to this subclause, the viewIdDepthTargetList is inferred to be identical to viewIdTargetList. viewIdDepthTargetList shall not be present as input if depthPresentFlagTarget is equal to 0.

The sub-bitstream is derived by applying the following operations in sequential order:

1. Let VOIDxList be empty and minVOIDx be the VOIDx value of the base view.
2. For each value of viewIdTarget included in viewIdTargetList, invoke the process specified in subclause I.8.5.1 for texture views with the viewIdTarget as input.
3. If depthPresentFlagTarget is equal to 1, for each value of viewIdTarget included in viewIdDepthTargetList, invoke the process specified in subclause I.8.5.1 for depth views with the viewIdTarget as input.
4. For each value of viewIdTarget included in viewIdTargetList, invoke the process specified in subclause I.8.5.2 for texture views with the value of viewIdTarget as input.
5. If depthPresentFlagTarget is equal to 1, for each value of viewIdTarget included in viewIdDepthTargetList, invoke the process specified in subclause I.8.5.2 for depth views with the viewIdTarget as input.
6. Mark all VCL NAL units of texture view components and filler data NAL units for which any of the following conditions are true as "to be removed from the bitstream":
 - priority_id is greater than pIdTarget,
 - temporal_id is greater than tIdTarget,
 - anchor_pic_flag is equal to 1, nal_unit_type is not equal to 21 and view_id is not marked as "required for anchor texture",
 - anchor_pic_flag is equal to 0, nal_unit_type is not equal to 21 and view_id is not marked as "required for non-anchor texture",
 - anchor_pic_flag is equal to 1, nal_unit_type is equal to 21 and view_id is not marked as "required for anchor depth",
 - anchor_pic_flag is equal to 0, nal_unit_type is equal to 21 and view_id is not marked as "required for non-anchor depth",
 - nal_unit_type is not equal to 21, nal_ref_idc is equal to 0 and inter_view_flag is equal to 0 and view_id is not equal to any value in the list viewIdTargetList.
 - nal_unit_type is equal to 21, nal_ref_idc is equal to 0 and inter_view_flag is equal to 0 and view_id is not equal to any value in the list viewIdDepthTargetList.
 - nal_unit_type is equal to 21 and depthPresentFlagTarget is equal to 0.
7. Remove all access units for which all VCL NAL units are marked as "to be removed from the bitstream".
8. Remove all VCL NAL units and filler data NAL units that are marked as "to be removed from the bitstream".
9. When VOIDxList contains only one value of VOIDx that is equal to minVOIDx, remove the following NAL units:

- all NAL units with `nal_unit_type` equal to 14 or 15,
- all NAL units with `nal_unit_type` equal to 6 in which the first SEI message has `payloadType` in the range of 36 to 44, inclusive.

NOTE 3 – When `VOIDxList` contains only one value of `VOIDx` equal to `minVOIDx`, the sub-bitstream contains only the base view or only a temporal subset of the base view.

10. Remove all NAL units with `nal_unit_type` equal to 6 in which the first SEI message has `payloadType` equal to 0 or 1, or the first SEI message has `payloadType` equal to 37 (MVC scalable nesting SEI message) and `operation_point_flag` in the first SEI message is equal to 1.

NOTE 4 – The buffering period SEI and picture timing SEI messages, when not nested or nested in the MVC scalable nesting SEI message, apply for a sub-bitstream obtained with the sub-bitstream extraction process of subclause H.8.5.3, which does not process NAL units of `nal_unit_type` equal to 21.

11. When `depthPresentFlagTarget` is equal to 0, the following applies in sequential order.

- Replace all NAL units with `nal_unit_type` equal to 6 in which `payloadType` indicates a MVCD texture scalable nesting SEI message with `sei_op_texture_only_flag` equal to 0 with a MVC scalable nesting SEI message with the same values of `num_view_components_op_minus1`, `sei_op_view_id[ i ]` and `sei_op_temporal_id` and the same nested SEI messages.
- Remove all NAL units with `nal_unit_type` equal to 6 in which `payloadType` indicates a MVCD texture scalable nesting SEI message.
- The following applies for each active texture MVCD sequence parameter set RBSP.
 - Replace `mvc_vui_parameters_extension()` syntax structure in an active texture MVCD sequence parameter set RBSPs with the `mvc_vui_parameters_extension()` syntax structure of the MVCD texture sub-bitstream VUI parameters extension, if both `mvc_vui_parameters_extension()` syntax structures apply to the same views.
 - Otherwise, remove `mvc_vui_parameters_extension()` syntax structure in an active texture MVCD sequence parameter set RBSP.
 - When `depthPresentFlagTarget` is equal to 0, remove all NAL units with `nal_unit_type` equal to 6 in which the first SEI message has `payloadType` in the range of 48 to 52, inclusive.

12. Let `maxTId` be the maximum `temporal_id` of all the remaining VCL NAL units. Remove all NAL units with `nal_unit_type` equal to 6 that only contain SEI messages that are part of an MVC scalable nesting SEI message or MVCD scalable nesting SEI message with any of the following properties:

- `operation_point_flag` is equal to 0 and `all_view_components_in_au_flag` is equal to 0 and none of `sei_view_id[ i ]` for all `i` in the range of 0 to `num_view_components_minus1`, inclusive, corresponds to a `VOIDx` value included in `VOIDxList`,
- `operation_point_flag` is equal to 1 and either `sei_op_temporal_id` is greater than `maxTId` or the list of `sei_op_view_id[ i ]` for all `i` in the range of 0 to `num_view_components_op_minus1`, inclusive, is not a subset of `viewIdTargetList` (i.e., it is not true that `sei_op_view_id[ i ]` for any `i` in the range of 0 to `num_view_components_op_minus1`, inclusive, is equal to a value in `viewIdTargetList`).

13. Remove each view scalability information SEI message and each operation point not present SEI message, when present.

14. When VOIdxList does not contain a value of VOIdx equal to minVOIdx, the view with VOIdx equal to the minimum VOIdx value included in VOIdxList is converted to the base view of the extracted sub-bitstream. An informative procedure that outlines key processing steps to create a base view is described in subclause I.8.5.6.

NOTE 5 – When VOIdxList does not contain a value of VOIdx equal to minVOIdx, the resulting sub-bitstream according to the operation steps 1-9 above does not contain a base view that conforms to one or more profiles specified in Annex A. In this case, by this operation step, the remaining view with the new minimum VOIdx value is converted to be the new base view that conforms to one or more profiles specified in Annex A and Annex H.

I.8.5.4 Specification of the base view bitstream

A bitstream that conforms to one or more profiles as specified in Annex I shall contain a base view bitstream that conforms to one or more of the profiles specified in Annex A. This base view bitstream is derived by invoking the sub-bitstream extraction process as specified in subclause I.8.5.3 with no input and the base view bitstream being the output.

NOTE – Although all multiview bitstreams that conform to one or more of the profiles specified in this annex contain a base view bitstream that conforms to one or more of the profiles specified in Annex A, the complete multiview bitstream (prior to operation of the base view extraction process specified in this subclause) may not conform to any profile specified in Annex A.

I.8.5.5 Specification of the stereoscopic texture bitstream

A bitstream that conforms to a profile as specified in Annex I shall contain at least one sub-bitstream that conforms to one or more of the profiles specified in Annex H with number of views equal to 2. This stereoscopic texture bitstream is derived by invoking the sub-bitstream extraction process as specified in subclause I.8.5.3 with depthPresentFlagTarget equal to 0 and viewIdTargetList containing the view_id values of the base view and a non-base view, the texture of which does not depend on any other non-base view for decoding.

I.9 Parsing process

The specifications in clause 9 apply.

I.10 Profiles and levels

The specifications in Annex H apply. Additional profiles and specific values of profile_idc are specified in the following.

The profiles that are specified in subclause I.10.1 are also referred to as the profiles specified in Annex I.

I.10.1 Profiles

All constraints for picture parameter sets that are specified in the following are constraints for picture parameter sets that become the active picture parameter set or an active view picture parameter set inside the bitstream. All constraints for MVCD sequence parameter sets that are specified in the following are constraints for MVCD sequence parameter sets that become the active MVCD sequence parameter set or an active view MVCD sequence parameter set inside the bitstream.

I.10.1.1 Multiview Depth High Profile

Bitstreams conforming to the Multiview Depth High profile shall obey the following constraints:

- The base view bitstream as specified in subclause I.8.5.4 shall obey all constraints of the High profile specified in subclause A.2.4 and all active sequence parameter sets shall fulfill one of the following conditions:
 - profile_idc is equal to 77 or constraint_set1_flag is equal to 1,
 - profile_idc is equal to 100.
- The sub-bitstream of stereoscopic texture bitstream as specified in subclause I.8.5.5 shall obey all constraints of the Stereo High profile specified in subclause H.10.2 and all active MVC sequence parameter sets shall fulfill one of the following conditions:

- profile_idc is equal to 128,
 - profile_idc is equal to 118 and constraint_set5_flag is equal to 1,
 - profile_idc is equal to 100,
 - profile_idc is equal to 77 or constraint_set1_flag is equal to 1.
 - Only I, P, and B slice types may be present.
 - NAL unit streams shall not contain nal_unit_type values in the range of 2 to 4, inclusive.
 - Arbitrary slice order is not allowed.
 - Picture parameter sets shall have num_slice_groups_minus1 equal to 0 only.
 - Picture parameter sets shall have redundant_pic_cnt_present_flag equal to 0 only.
 - When frame_mbs_only_flag is equal to 1 in an active sequence parameter set for a texture view, frame_mbs_only_flag shall be equal to 1 in the active sequence parameter set for the depth view having the same view_id.
 - When frame_mbs_only_flag is equal to 0 in an active sequence parameter set for a depth view, mb_adaptive_frame_field_flag shall be equal to 0.
 - MVCD sequence parameter sets for the depth views shall have chroma_format_idc equal to 0 only.
 - MVCD sequence parameter sets shall have bit_depth_luma_minus8 equal to 0 only.
 - MVCD sequence parameter sets shall have bit_depth_chroma_minus8 equal to 0 only.
 - MVCD sequence parameter sets shall have qprime_y_zero_transform_bypass_flag equal to 0 only.
 - For each access unit, the value of level_idc for all active view MVCD sequence parameter set RBSPs shall be the same as the value of level_idc for the active MVCD sequence parameter set RBSP.
 - The level constraints specified for the Multiview Depth High profile in subclause I.10.2 shall be fulfilled.
- Conformance of a bitstream to the Multiview Depth High profile is indicated by profile_idc being equal to 138.

Decoders conforming to the Multiview Depth High profile at a specific level shall be capable of decoding all bitstreams in which both of the following conditions are true:

- a) All active MVCD sequence parameter sets have one or more of the following conditions fulfilled:
 - profile_idc is equal to 138,
 - profile_idc is equal to 128,
 - profile_idc is equal to 118 and constraint_set5_flag is equal to 1,
 - profile_idc is equal to 100,
 - profile_idc is equal to 77 or constraint_set1_flag is equal to 1.
- b) All active MVCD sequence parameter sets have one or more of the following conditions fulfilled:
 - level_idc or (level_idc and constraint_set3_flag) represent a level less than or equal to the specific level,
 - level_idc[i] or (level_idc[i] and constraint_set3_flag) represent a level less than or equal to the specific level.

I.10.2 Levels

The following is specified for expressing the constraints in this clause:

- Let access unit n be the n-th access unit in decoding order with the first access unit being access unit 0.
- Let picture n be the primary coded picture or the corresponding decoded picture of access unit n.

Let the variable fR be derived as follows:

- If picture n is a frame, fR is set equal to $1 \div 172$.
- Otherwise (picture n is a field), fR is set equal to $1 \div (172 * 2)$.

The value of $mvcdScaleFactor$ is set equal to 2.

The value of $mvcdScaleFactor$ is set equal to 2.5.

The value of $NumViews$ indicates the number of views, including texture views and depth views, which are required for decoding the target output views corresponding to the j -th operation point for $level_idc[i]$ as signalled in the subset sequence parameter set, and is set equal to $applicable_op_num_depth_views_minus1[i][j] + applicable_op_num_depth_views_minus1[i][j] + 2$.

The value of $PicWidthInMbs$ and $FrameHeightInMbs$ refer to the width and height of each view component, while the value of $TotalPicSizeInMbs$ indicates the total number of macroblocks in the texture view components and depth view components of a picture.

I.10.2.1 Level limits common to Multiview Depth High profiles

Bitstreams conforming to the Multiview Depth High profile at a specified level shall obey the following constraints:

- a) The nominal removal time of access unit n (with $n > 0$) from the CPB as specified in clause C.1.2, satisfies the constraint that $t_{r,n}(n) - t_r(n-1)$ is greater than or equal to $\text{Max}(TotalPicSizeInMbs \div (mvcdScaleFactor * MaxMBPS), fR)$, where $MaxMBPS$ is the value specified in Table A-1 that applies to picture $n-1$, and $TotalPicSizeInMbs$ is the total number of macroblocks in the texture view components and depth view components of picture $n-1$.
- b) The difference between consecutive output times of pictures from the DPB as specified in clause C.2.2, satisfies the constraint that $\Delta t_{o,dpb}(n) \geq \text{Max}(TotalPicSizeInMbs \div (mvcdScaleFactor * MaxMBPS), fR)$, where $MaxMBPS$ is the value specified in Table A-1 for picture n , and $TotalPicSizeInMbs$ is the total number of macroblocks in the texture view components and depth view components of picture n , provided that picture n is a picture that is output and is not the last picture of the bitstream that is output.
- c) $PicWidthInMbs * FrameHeightInMbs \leq MaxFS$, where $MaxFS$ is specified in Table A-1.
- d) $PicWidthInMbs \leq \text{Sqrt}(MaxFS * 8)$, where $MaxFS$ is specified in Table A-1.
- e) $FrameHeightInMbs \leq \text{Sqrt}(MaxFS * 8)$, where $MaxFS$ is specified in Table A-1.
- f) $max_dec_frame_buffering \leq \text{MaxDpbFrames}$, where $MaxDpbFrames$ is equal to $\text{Min}(mvcdScaleFactor * MaxDpbMbs / (TotalPicSizeInMbs / NumViews), \text{Max}(1, \text{Ceil}(\log_2(NumViews))) * 16)$ and $MaxDpbMbs$ is specified in Table A-1.
- g) The vertical motion vector component range does not exceed $MaxVmvR$ in units of luma frame samples, where $MaxVmvR$ is specified in Table A-1.
- h) The horizontal motion vector range does not exceed the range of -2048 to 2047.75 , inclusive, in units of luma samples.
- i) Let $setOf2Mb$ be the set of unsorted pairs of macroblocks that contains the unsorted pairs of macroblocks (mbA, mbB) of a coded video sequence for which any of the following conditions are true:
 - mbA and mbB are macroblocks that belong to the same slice and are consecutive in decoding order,
 - $separate_colour_plane_flag$ is equal to 0, mbA is the last macroblock (in decoding order) of a slice, and mbB is the first macroblock (in decoding order) of the next slice in decoding order,
 - $separate_colour_plane_flag$ is equal to 1, mbA is the last macroblock (in decoding order) of a slice with a particular value of $colour_plane_id$, and mbB is the first macroblock (in decoding order) of the next slice with the same value of $colour_plane_id$ in decoding order.

NOTE 1 – In the two above conditions, the macroblocks mbA and mbB can belong to different pictures.

For each unsorted pair of macroblocks (mbA, mbB) of the set setOf2Mb, the total number of motion vectors (given by the sum of the number of motion vectors for macroblock mbA and the number of motion vectors for macroblock mbB) does not exceed MaxMvsPer2Mb, where MaxMvsPer2Mb is specified in Table A-1. The number of motion vectors for each macroblock is the value of the variable MvCnt after the completion of the intra or inter prediction process for the macroblock.

NOTE 2 – When separate_colour_plane_flag is equal to 0, the constraint specifies that the total number of motion vectors for two consecutive macroblocks in decoding order must not exceed MaxMvsPer2Mb. When separate_colour_plane_flag is equal to 1, the constraint specifies that the total number of motion vectors for two consecutive macroblocks with the same value of colour_plane_id in decoding order must not exceed MaxMvsPer2Mb. For macroblocks that are consecutive in decoding order but are associated with a different value of colour_plane_id, no constraint for the total number of motion vectors is specified.

- j) The number of bits of macroblock_layer() data for any macroblock is not greater than $128 + \text{RawMbBits}$. Depending on entropy_coding_mode_flag, the bits of macroblock_layer() data are counted as follows:
 - If entropy_coding_mode_flag is equal to 0, the number of bits of macroblock_layer() data is given by the number of bits in the macroblock_layer() syntax structure for a macroblock.
 - Otherwise (entropy_coding_mode_flag is equal to 1), the number of bits of macroblock_layer() data for a macroblock is given by the number of times read_bits(1) is called in clauses 9.3.3.2.2 and 9.3.3.2.3 when parsing the macroblock_layer() associated with the macroblock.
- k) The removal time of access unit 0 shall satisfy the constraint that the number of slices in picture 0 is less than or equal to $\text{mvcdScaleFactor} * (\text{Max}(\text{PicSizeInMbs}, \text{fR} * \text{MaxMBPS}) + \text{MaxMBPS} * (\text{t}_r(0) - \text{t}_{r,n}(0))) \div \text{SliceRate}$, where MaxMBPS and SliceRate are the values specified in Tables A-1 and A-4, respectively, that apply to picture 0 and PicSizeInMbs is the number of macroblocks in a single texture view component of picture 0.
- l) The removal time of access unit 0 shall satisfy the constraint that the number of slices in each view component of picture 0 is less than or equal to $(\text{Max}(\text{PicSizeInMbs}, \text{fR} * \text{MaxMBPS}) + \text{MaxMBPS} * (\text{t}_r(0) - \text{t}_{r,n}(0))) \div \text{SliceRate}$, where MaxMBPS and SliceRate are the values specified in Tables A-1 and A-4, respectively, that apply to picture 0 and PicSizeInMbs is the number of macroblocks in a single view component of picture 0.
- m) The difference between consecutive removal times of access units n and n – 1 with n > 0 shall satisfy the constraint that the number of slices in picture n is less than or equal to $\text{mvcdScaleFactor} * \text{MaxMBPS} * (\text{t}_r(n) - \text{t}_r(n-1)) \div \text{SliceRate}$, where SliceRate is the value specified in Table A-4 that applies to picture n.
- n) The difference between consecutive removal times of access units n and n – 1 with n > 0 shall satisfy the constraint that the number of slices in each view component of picture n is less than or equal to $\text{MaxMBPS} * (\text{t}_r(n) - \text{t}_r(n-1)) \div \text{SliceRate}$, where SliceRate is the value specified in Table A-4 that applies to picture n.
- o) MVCD sequence parameter sets shall have direct_8x8_inference_flag equal to 1 for the levels specified in Table A-4.
- p) The value of sub_mb_type[mbPartIdx] with mbPartIdx = 0..3 in B macroblocks with mb_type equal to B_8x8 shall not be equal to B_Bi_8x4, B_Bi_4x8, or B_Bi_4x4 for the levels in which MinLumaBiPredSize is shown as 8x8 in Table A-4.
- q) For the VCL HRD parameters, $\text{BitRate}[\text{SchedSelIdx}] \leq \text{cpbBrVclFactor} * \text{MaxBR}$ and $\text{CpbSize}[\text{SchedSelIdx}] \leq \text{cpbBrVclFactor} * \text{MaxCPB}$ for at least one value of SchedSelIdx, where cpbBrVclFactor is equal to 1250. With vui_mv_c_vcl_hrd_parameters_present_flag[i] being the syntax element, in the MVCD VUI parameters extension of the active MVCD sequence parameter set, that is associated with the VCL HRD parameters that are used for conformance checking (as specified in Annex C), BitRate[SchedSelIdx] and CpbSize[SchedSelIdx] are given as follows:

- If `vui_mvc_vcl_hrd_parameters_present_flag` equal to 1, `BitRate[ SchedSelIdx ]` and `CpbSize[ SchedSelIdx ]` are given by Equations E-37 and E-38, respectively, using the syntax elements of the `hrd_parameters()` syntax structure that immediately follows `vui_mvc_vcl_hrd_parameters_present_flag`.
- Otherwise (`vui_mvc_vcl_hrd_parameters_present_flag` equal to 0), `BitRate[ SchedSelIdx ]` and `CpbSize[ SchedSelIdx ]` are inferred as specified in clause E.2.2 for VCL HRD parameters.

MaxBR and MaxCPB are specified in Table A-1 in units of `cpbBrVclFactor` bits/s and `cpbBrVclFactor` bits, respectively. The bitstream shall satisfy these conditions for at least one value of `SchedSelIdx` in the range 0 to `cpb_cnt_minus1`, inclusive.

- r) For the NAL HRD parameters, `BitRate[ SchedSelIdx ]` $\leq$ `cpbBrNalFactor` * MaxBR and `CpbSize[ SchedSelIdx ]` $\leq$ `cpbBrNalFactor` * MaxCPB for at least one value of `SchedSelIdx`, where `cpbBrNalFactor` is equal to 1500. With `vui_mvc_nal_hrd_parameters_present_flag[ i ]` being the syntax element, in the MVCD VUI parameters extension of the active MVCD sequence parameter set, that is associated with the NAL HRD parameters that are used for conformance checking (as specified in Annex C), `BitRate[ SchedSelIdx ]` and `CpbSize[ SchedSelIdx ]` are given as follows:

- If `vui_mvc_nal_hrd_parameters_present_flag` equal to 1, `BitRate[ SchedSelIdx ]` and `CpbSize[ SchedSelIdx ]` are given by Equations E-37 and E-38, respectively, using the syntax elements of the `hrd_parameters()` syntax structure that immediately follows `vui_mvc_nal_hrd_parameters_present_flag`.
- Otherwise (`vui_mvc_nal_hrd_parameters_present_flag` equal to 0), `BitRate[ SchedSelIdx ]` and `CpbSize[ SchedSelIdx ]` are inferred as specified in clause E.2.2 for NAL HRD parameters.

MaxBR and MaxCPB are specified in Table A-1 in units of `cpbBrNalFactor` bits/s and `cpbBrNalFactor` bits, respectively. The bitstream shall satisfy these conditions for at least one value of `SchedSelIdx` in the range 0 to `cpb_cnt_minus1`, inclusive.

- s) The sum of the `NumBytesInNALunit` variables for access unit 0 is less than or equal to $384 * mvcdScaleFactor * (\text{Max}(\text{PicSizeInMbs}, fR * \text{MaxMBPS}) + \text{MaxMBPS} * (t_r(0) - t_{r,n}(0))) \div \text{MinCR}$, where `MaxMBPS` and `MinCR` are the values specified in Table A-1 that apply to picture 0 and `PicSizeInMbs` is the number of macroblocks in a single texture view component of picture 0.
- t) The sum of the `NumBytesInNALunit` variables for the VCL NAL units of each view component of access unit 0 is less than or equal to $384 * (\text{Max}(\text{PicSizeInMbs}, fR * \text{MaxMBPS}) + \text{MaxMBPS} * (t_r(0) - t_{r,n}(0))) \div \text{MinCR}$, where `MaxMBPS` and `MinCR` are the values specified in Table A-1 that apply to picture 0 and `PicSizeInMbs` is the number of macroblocks in a single view component of picture 0.
- u) The sum of the `NumBytesInNALunit` variables for access unit `n` with `n > 0` is less than or equal to $384 * mvcdScaleFactor * \text{MaxMBPS} * (t_r(n) - t_r(n-1)) \div \text{MinCR}$, where `MaxMBPS` and `MinCR` are the values specified in Table A-1 that apply to picture `n`.
- v) The sum of the `NumBytesInNALunit` variables for the VCL NAL units of each view component of access unit `n` with `n > 0` is less than or equal to $384 * \text{MaxMBPS} * (t_r(n) - t_r(n-1)) \div \text{MinCR}$, where `MaxMBPS` and `MinCR` are the values specified in Table A-1 that apply to picture `n`.
- w) When `PicSizeInMbs` is greater than 1620, the number of macroblocks in any coded slice shall not exceed `MaxFS / 4`, where `MaxFS` is specified in Table A-1.
- x) `max_num_ref_frames` shall be less than or equal to `MaxDpbFrames / mvcdScaleFactor` for each texture view component, where `MaxDpbFrames` is specified in item f).
- y) MVCD sequence parameter sets shall have `frame_mbs_only_flag` equal to 1 for the levels specified in Table A-4.

Table A-1 specifies the limits for each level. A definition of all levels identified in the "Level number" column of Table A-1 is specified for the Multiview Depth High profile. Table A-4 specifies limits for each level that are specific to bitstreams conforming to the Multiview Depth High profile. Each entry in Tables A-1 and A-4 indicates, for the level corresponding to the row of the table, the absence or value of a limit that is imposed by the variable corresponding to the column of the table, as follows:

- If the table entry is marked as "-", no limit is imposed by the value of the variable as a requirement of bitstream conformance to the profile at the specified level.
- Otherwise, the table entry specifies the value of the variable for the associated limit that is imposed as a requirement of bitstream conformance to the profile at the specified level.

For coded video sequences conforming to the Multiview Depth High profile, the level_idc value is specified as follows:

- If level_idc is not equal to 0, level_idc indicates the level that applies to the coded video sequence operating with all the views being target output views.

NOTE 3 – A level_idc value that is not equal to zero may indicate a higher level than necessary to decode the coded video sequence operating with all the views being target output views. This may occur when a subset of views or temporal subsets are removed from a coded video sequence according to the sub-bitstream extraction process specified in clause I.8.5.3, and the level_idc value is not updated accordingly.

- Otherwise (level_idc is equal to 0), the level that applies to the coded video sequence operating with all the views being target output views is unspecified.

NOTE 4 – When profile_idc is equal to 118 or 128 and level_idc is equal to 0, there may exist a level indicated by level_idc[i] that is applicable to the coded video sequence operating with all the views being target output views. This may occur when a subset of views or temporal subsets are removed from a coded video sequence according to the sub-bitstream extraction process specified in clause I.8.5.3, and a particular value of level_idc[i] corresponds to the resulting coded video sequence.

In bitstreams conforming to the Multiview Depth High profile, the conformance of the bitstream to a specified level is indicated by the syntax element level_idc or level_idc[i] as follows:

- If level_idc or level_idc[i] is equal to 9, the indicated level is level 1b.
- Otherwise (level_idc or level_idc[i] is not equal to 9), level_idc or level_idc[i] is equal to a value of ten times the level number (of the indicated level) specified in Table A-1.

I.10.2.2 Profile specific level limits

- a) In bitstreams conforming to the Multiview Depth High profile, MVCD sequence parameter sets shall have frame_mbs_only_flag equal to 1 for the levels specified in Table A-4.

I.11 Byte stream format

The specifications in Annex B apply.

I.12 MVCD hypothetical reference decoder

The specifications in Annex C apply with substituting MVCD sequence parameter set for MVC sequence parameter set.

I.13 MVCD SEI messages

The specifications in Annex D together with the extensions and modifications specified in this subclause apply.

I.13.1 SEI message syntax

I.13.1.1 MVCD view scalability information SEI message syntax

	C	Descriptor
mvcd_view_scalability_info(payloadSize) {		
num_operation_points_minus1	5	ue(v)
for(i = 0; i <= num_operation_points_minus1; i++) {		
operation_point_id[i]	5	ue(v)
priority_id[i]	5	u(5)
temporal_id[i]	5	u(3)
num_target_output_views_minus1[i]	5	ue(v)
for(j = 0; j <= num_target_output_views_minus1[i]; j++) {		
view_id[i][j]	5	ue(v)
mvcd_op_view_info()		
}		
profile_level_info_present_flag[i]	5	u(1)
bitrate_info_present_flag[i]	5	u(1)
frm_rate_info_present_flag[i]	5	u(1)
if(!num_target_output_views_minus1[i])		
view_dependency_info_present_flag[i]	5	u(1)
parameter_sets_info_present_flag[i]	5	u(1)
bitstream_restriction_info_present_flag[i]	5	u(1)
if(profile_level_info_present_flag[i])		
op_profile_level_idc[i]	5	u(24)
if(bitrate_info_present_flag[i]) {		
avg_bitrate[i]	5	u(16)
max_bitrate[i]	5	u(16)
max_bitrate_calc_window[i]	5	u(16)
}		
if(frm_rate_info_present_flag[i]) {		
constant_frm_rate_idc[i]	5	u(2)
avg_frm_rate[i]	5	u(16)
}		
if(view_dependency_info_present_flag[i]) {		
num_directly_dependent_views[i]	5	ue(v)
for(j = 0; j < num_directly_dependent_views[i]; j++) {		
directly_dependent_view_id[i][j]	5	ue(v)
mvcd_op_view_info()		
}		
} else		
view_dependency_info_src_op_id[i]	5	ue(v)
if(parameter_sets_info_present_flag[i]) {		
num_seq_parameter_set_minus1[i]	5	ue(v)
for(j = 0; j <= num_seq_parameter_set_minus1[i]; j++)		
seq_parameter_set_id_delta[i][j]	5	ue(v)
num_subset_seq_parameter_set_minus1[i]	5	ue(v)
for(j = 0; j <= num_subset_seq_parameter_set_minus1[i]; j++)		
subset_seq_parameter_set_id_delta[i][j]	5	ue(v)
num_pic_parameter_set_minus1[i]	5	ue(v)

for(j = 0; j <= num_init_pic_parameter_set_minus1[i]; j++)		
pic_parameter_set_id_delta[i][j]	5	ue(v)
} else		
parameter_sets_info_src_op_id[i]	5	ue(v)
if(bitstream_restriction_info_present_flag[i]) {		
motion_vectors_over_pic_boundaries_flag[i]	5	u(1)
max_bytes_per_pic_denom[i]	5	ue(v)
max_bits_per_mb_denom[i]	5	ue(v)
log2_max_mv_length_horizontal[i]	5	ue(v)
log2_max_mv_length_vertical[i]	5	ue(v)
num_reorder_frames[i]	5	ue(v)
max_dec_frame_buffering[i]	5	ue(v)
}		
}		
}		

I.13.1.1.1 MVCD operation point view information syntax

mvcd_op_view_info() {	C	Descriptor
view_info_depth_view_present_flag	5	u(1)
if(view_info_depth_view_present_flag)		
reserved_depth_view_confirmation_flag	5	u(1)
view_info_texture_view_present_flag	5	u(1)
if(view_info_texture_view_present_flag)		
reserved_texture_view_confirmation_flag	5	u(1)
}		

I.13.1.1.2 MVCD scalable nesting SEI message syntax

mvcd_scalable_nesting(payloadSize) {	C	Descriptor
operation_point_flag	5	u(1)
if(!operation_point_flag) {		
all_view_components_in_au_flag	5	u(1)
if(!all_view_components_in_au_flag) {		
num_view_components_minus1	5	ue(v)
for(i = 0; i <= num_view_components_minus1; i++) {		
sei_view_id[i]	5	u(10)
sei_view_applicability_flag[i]	5	u(1)
}		
}		
} else {		
sei_op_texture_only_flag	5	u(1)
num_view_components_op_minus1	5	ue(v)
for(i = 0; i <= num_view_components_op_minus1; i++) {		
sei_op_view_id[i]	5	u(10)
if(!sei_op_texture_only_flag) {		
sei_op_depth_flag[i]		
sei_op_texture_flag[i]		

}		
}		
sei_op_temporal_id	5	u(3)
}		
while(!byte_aligned())		
sei_nesting_zero_bit /* equal to 0 */	5	f(1)
sei_message()	5	
}		

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC 14496-10:2012/Amd 2:2013

I.13.1.3 Depth representation information SEI message syntax

depth_representation_info(payloadSize) {	C	Descriptor
all_views_equal_flag	5	u(1)
if(all_views_equal_flag == 0) {		
num_views_minus1	5	ue(v)
numViews = num_views_minus1 + 1		
} else		
numViews = 1		
z_near_flag	5	u(1)
z_far_flag	5	u(1)
if(z_near_flag z_far_flag) {		
z_axis_equal_flag	5	u(1)
if(z_axis_equal_flag)		
common_z_axis_reference_view	5	ue(v)
}		
d_min_flag	5	u(1)
d_max_flag	5	u(1)
depth_representation_type	5	ue(v)
for(i = 0; i < numViews; i++) {		
depth_info_view_id[i]	5	ue(v)
if((z_near_flag z_far_flag) && (z_axis_equal_flag == 0))		
z_axis_reference_view[i]	5	ue(v)
if(d_min_flag d_max_flag)		
disparity_reference_view[i]	5	ue(v)
if(z_near_flag)		
depth_representation_sei_element(ZNearSign, ZNearExp, ZNearMantissa, ZNearManLen)		
if(z_far_flag)		
depth_representation_sei_element(ZFarSign, ZFarExp, ZFarMantissa, ZFarManLen)		
if(d_min_flag)		
depth_representation_sei_element(DMinSign, DMinExp, DMinMantissa, DMinManLen)		
if(d_max_flag)		
depth_representation_sei_element(DMaxSign, DMaxExp, DMaxMantissa, DMaxManLen)		
}		
if(depth_representation_type == 3) {		
depth_nonlinear_representation_num_minus1	5	ue(v)
for(i = 1; i <= depth_nonlinear_representation_num_minus1 + 1; i++)		
depth_nonlinear_representation_model[i]	5	ue(v)
}		
}		

I.13.1.4 Depth representation SEI element syntax

depth_representation_sei_element(OutSign, OutExp, OutMantissa, OutManLen) {	C	Descriptor
da_sign_flag	5	u(1)
da_exponent	5	u(7)
da_mantissa_len_minus1	5	u(5)
da_mantissa	5	u(v)
}		

I.13.1.5 3D reference displays information SEI message syntax

three_dimensional_reference_displays_info(payloadSize) {	C	Descriptor
prec_ref_baseline	5	ue(v)
prec_ref_display_width	5	ue(v)
ref_viewing_distance_flag	5	u(1)
if(ref_viewing_distance_flag)		
prec_ref_viewing_dist	5	ue(v)
num_ref_displays_minus1	5	ue(v)
numRefDisplays = num_ref_displays_minus1 + 1		
for(i = 0; i < numRefDisplays; i++) {		
exponent_ref_baseline[i]	5	u(6)
mantissa_ref_baseline[i]	5	u(v)
exponent_ref_display_width[i]	5	u(6)
mantissa_ref_display_width[i]	5	u(v)
if(ref_viewing_distance_flag) {		
exponent_ref_viewing_distance[i]	5	u(6)
mantissa_ref_viewing_distance[i]	5	u(v)
}		
additional_shift_present_flag[i]	5	u(1)
if(additional_shift_present[i])		
num_sample_shift_plus512[i]	5	u(10)
}		
three_dimensional_reference_displays_extension_flag	5	u(1)
}		

I.13.1.6 Depth timing SEI message syntax

depth_timing(payloadSize) {	C	Descriptor
per_view_depth_timing_flag	5	u(1)
if(per_view_depth_timing_flag)		
for(i = 0; i < NumDepthViews; i++)		
depth_timing_offset()		
else		
depth_timing_offset()		
}		

I.13.1.6.1 Depth timing offset syntax

depth_timing_offset() {	C	Descriptor
--------------------------	---	------------