

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
26 March 2009 (26.03.2009)

PCT

(10) International Publication Number
WO 2009/037672 A2

(51) International Patent Classification: **Not classified**

(21) International Application Number:
PCT/IB2008/053823

(22) International Filing Date:
20 September 2008 (20.09.2008)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
60/974,401 21 September 2007 (21.09.2007) US
60/991,640 30 November 2007 (30.11.2007) US

(71) Applicant (for all designated States except US):
NOKIA CORPORATION [FI/FI]; Keilalahdentie 4,
FI-FIN-02150 Espoo (FI).

(71) Applicant (for LC only): **NOKIA INC** [US/US]; 6000
Connection Drive, Irving, Texas 75039 (US).

(72) Inventors; and

(75) Inventors/Applicants (for US only): **OJALA, Pasi**

[FI/FI]; Neidonkalliontie 9, FI-02400 Kirkkonummi (FI).
LAKANIEMI, Ari [FI/FI]; Laajalahdentie 24 A 10,
FI-00330 Helsinki (FI).

(74) Agent: **ALBERT, JR., G. Peter**; Foley & Lardner, LLP,
P.O. Box 80278, San Diego, CA 92130 (US).

(81) Designated States (unless otherwise indicated, for every
kind of national protection available): AE, AG, AL, AM,
AO, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA,
CH, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE,
EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID,
IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK,
LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW,
MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL, PT,
RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TJ,
TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM,
ZW.

(84) Designated States (unless otherwise indicated, for every
kind of regional protection available): ARIPO (BW, GH,
GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM,
ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM),

[Continued on next page]

(54) Title: **SCALABLE ERROR DETECTION AND CROSS-SESSION TIMING SYNCHRONIZATION FOR
PACKET-SWITCHED TRANSMISSION**




FIG. 1

(57) Abstract: A system and method for providing improved scalable error detection and cross- timing synchronization for packet-switched transmission. In one embodiment, checksum error detection is applied for the core layer and for enhancement layers of the scalable payload in such a way that dropping one or several enhancement layers from the payload does not change the value of the checksum. Only one checksum is transmitted, e.g., in the payload or in the header of the lower-layer protocol. The transmitter modifies the encoded bit stream in such a manner that the entity in the network deploying the scalable payload and removing layers from the packet does not need to recalculate the checksum placed in the payload or packet header, even when the payload size is changed. A prefix/tail bit field is added in the beginning/end of each enhancement layer to make the checksum check match with the common checksum. In another embodiment, the receiver may check the correctness of each received layer simultaneously and, if desired, utilize data only from correctly received layers.



European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MT, NL, NO, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

— *without international search report and to be republished upon receipt of that report*

SCALABLE ERROR DETECTION AND CROSS-SESSION TIMING SYNCHRONIZATION FOR PACKET-SWITCHED TRANSMISSION

FIELD OF THE INVENTION

[0001] The present invention relates generally to packet-switched networks. More particularly, the present invention relates to scalable real-time media, media adaptation, the synchronization of Real-time Transport Protocol (RTP) sessions, error resiliency and transmission over packet switched networks, such as in Voice over IP (VoIP) and video conferencing.

BACKGROUND OF THE INVENTION

[0002] This section is intended to provide a background or context to the invention that is recited in the claims. The description herein may include concepts that could be pursued, but are not necessarily ones that have been previously conceived or pursued. Therefore, unless otherwise indicated herein, what is described in this section is not prior art to the description and claims in this application and is not admitted to be prior art by inclusion in this section.

[0003] Packet switched network transmission protocols typically employed for VoIP comprise the RTP encapsulated in the User Datagram Protocol (UDP), further encapsulated into the Internet Protocol (IP). The checksums employed in the UDP and IP result in discarding all the packets in which the receiver detects bit errors. In other words, the protocol stack in the receiver does not convey any distorted packets to the application layer. Therefore, when IP packets are transmitted over an error prone radio link or over any media introducing transmission errors, the application layer is likely to face packet losses. Conversely, none of the packets reaching the application layer contain any residual bit errors. Due to this phenomenon, an error

concealment algorithm is not able to utilize partially correct frames, as can be done, for example, in a circuit switched GSM telephone service. Instead, the erroneous frame needs to be completely replaced. This is likely to make the error concealment process less effective than the approach that is used in circuit-switched services.

[0004] Another aspect of packet-switched networks involves media level scalability. This scalability may be deployed in the transmission level, e.g., for controlling the network capacity or shaping a multicast media stream to facilitate operation with participants behind access links of different bandwidths. In the application level, the scalability can be utilized for controlling, e.g., computational complexity, encoding delay, or desired quality level.

[0005] The scalable media data comprises a core layer, which is always needed to enable reconstruction in the receiving end, and one or several enhancement layers that can be used to provide added value to the reconstructed media (e.g., improved media quality). It should be noted that, while in some scenarios the scalability can be applied in the transmitting end-point, there are also operating scenarios where it makes more sense to permit an intermediate network element to perform the scaling. The enhancement layers can be transmitted either together with the core layer data or in separate packets. While transmitting the enhancement layers in separate packets from the core layer data makes the scaling operation more straightforward since the scaling operation can be performed by dropping full packets, it provides some challenges, e.g. in the session setup and control procedures. On the other hand, the transmitting of the enhancement layers in the same packet together with the core layer data is more effective in terms of transmission bandwidth usage and also enables a simpler session setup.

[0006] Network congestion and limited link bandwidth are examples of reasons that may necessitate the removal of a portion of the scalable content from an IP packet. Furthermore, all receivers may not be capable of receiving or consuming full bit rate content. Hence, a network element controlling the transmission link may remove some of the higher layers of content.

[0007] When data from enhancement layers is transmitted in the same packet with core layer data, a scaling operation implies the modification of packet contents. The modification of an IP packet requires always opening and repacketizing the payload. Although the receiver and sender information do not change, the packet size and corresponding header information may be modified. For example, any error detection checksums need to be recalculated when the content of the payload is changed.

[0008] Media data that can be divided into two or more classes of importance can benefit from an approach where the error detection is performed separately for different classes of data. In the event that there are errors only in the “less important” parts of the data, many applications are still able to make use of the error-free part of the received data. A scalable media packet carrying the enhancement layers together with the core data clearly forms a similar case by its design; possible errors in the enhancement layers do not affect the core layer, and therefore the core data should be made available for the application also in case some of the enhancement layers are corrupted. A further benefit for the scalable media arises from the fact that separate error detection checksums for each of the layers facilitates simple scaling functionality. Furthermore, different forward error correction (FEC) arrangements can also be used for core and enhancement layers.

[0009] Various methods have been introduced to handle packet loss conditions. Some methods, such as partial checksum methods, involving handling only portions of the payload data. In particular, when a UDP-Lite protocol is at issue, a partial checksum could be utilized for unequal error detection. With scalable content, a problem arises involving the need for recalculating a partial checksum when the payload size is modified.

[0010] Typically, in circuit switched systems, the most sensitive bits of a speech codec are protected with stronger FEC arrangements compared to the least sensitive bits. A partial error detection code, such as cyclic redundancy check (CRC), can be used to classify the whole frame as lost when the most sensitive bits contain errors. A similar method can be used in packet switched networks.

[0011] Unequal FEC protection methods for speech frames are based on the ordering of parameters in descending order of priority. For example, the scalable codec bit stream can be classified to various classes according to the scalability layers and the importance of them. Each priority class can then be protected with an unequal error correction capability. Such an unequal protection arrangement requires two features for receiver implementations that are not typically standardized. First, the protocol stack, and especially the UDP and IP protocol implementations, should be modified to pass partially incorrect packet payloads to protocol stack layers above. Alternatively, protocols allowing erroneous payloads to be passed to the application, such as UDP-Lite, can be used.

[0012] An advantage of unequal error detection involves permitting increased capacity, particularly in wireless links when the errors in less sensitive bits do not cause the dropping of the whole packet. The application layer still copes with payloads containing errors in the least sensitive content.

[0013] One method for error detection involves having a single checksum covering the entire payload data. This is the approach used, for example, in the UDP protocol. This mode of operation does not enable UED at all and, in the case of scalable content, the removal of one or more enhancement layers from the payload requires re-computation of the checksum. While this is a fully working solution, the computational complexity of the additional CRC computation each time a packet is being scaled may not be a feasible approach in certain scenarios. On the other hand, only one checksum is needed to cover the whole payload in this approach, thereby providing an efficient usage of transmission bandwidth. Additionally, in this arrangement, the end-points only have to compute/verify a single checksum.

[0014] An enhancement to the single-checksum approach discussed above is to have a single checksum that covers only selected portions of the payload data, thereby leaving the rest of the payload data uncovered. This approach enables a simple UED, providing a checksum-covered part and uncovered parts of the payload data. Furthermore, a simple two-level scalability can be supported by having the core layer as the checksum-covered part and the enhancement layer as the uncovered part. In

this case, achieving scalability by dropping the enhancement layer would not require changing the checksum unless the change of the payload affects the value of the checksum. This kind of functionality is employed, e.g. in UDP-Lite by providing the ability to permit the checksum to cover only the selected part of the data at the beginning of the payload data. Making use of this UED-enabling functionality requires that the payload data must be arranged in a suitable way (i.e., the most important data should appear at the beginning of the payload) and the application can benefit by letting the errors in uncovered part go unnoticed. An example of such a payload is the RTP payload for the AMR and AMR-WB speech codecs, which provides an option for including a CRC bit field in the payload header. The CRC check is performed over the most sensitive speech bits, namely the class A bits. In this scenario, the goal is to discard the frame only if there are bit errors in the most sensitive bits. On the other hand, errors in the remaining portion of the frame can be tolerated. Currently, unequal error detection can be utilized in full only with the UDP-Lite and DCCP protocols. These protocols enable the utilization of functionality, which allows for the passing of partially corrupted packets up to the application layer.

[0015] In a variation of the above method, the UDP-Lite checksum, for example, can be used to cover part of the payload, with a separate payload-internal checksum being used to cover the parts of the payload not covered by the UDP-Lite checksum. This approach provides full checksum coverage for the payload, which enables also discarding the less important portion of the payload data in case errors are detected. This can be particularly useful in some applications that cannot tolerate any (undetected) errors.

[0016] A further option for a more flexible payload design involves dividing the data into several subsections and to have separate checksums for each subsection. This system enables both flexible UED and flexible scalability. When each of the layers is covered by a separate checksum, the dropping of any of the layers (either due to a detected error or due to a scaling operation) can be performed without the need to re-compute the checksum(s). However, a drawback to this system is that the use of

several checksums requires sending more data compared to previous approaches with only a single checksum. Furthermore, in this arrangement, several checksums need to be separately computed at the sender and verified at the receiving end. On the other hand, this approach enables an increased robustness to errors and flexible scalability, without the need to re-compute checksum(s) in an intermediate element that performs the scaling.

[0017] IP transport mechanisms provide tools for FEC packets. For example, the Internet Engineering Task Force (IETF) Request for Comments (RFC) 2733, which can be found at www.ietf.org/rfc/rfc2733.txt, provides a generic mechanism for transporting XOR-based forward error correction data within a separate RTP session. The payload header of FEC packets contains a bit mask identifying the packet payloads over which the bit-wise XOR operation is calculated and a few fields for RTP header recovery of the protected packets. One XOR FEC packet enables recovery of one lost source packet.

[0018] Work is being conducted to replace IETF RFC 2733 with similar RTP payload format for XOR-based FEC protection also including the capability of uneven levels of protection, herein referred to as the ULP Internet Draft (A.H. Li, "RTP payload format for generic forward error correction," Internet Engineering Task Force Internet Draft draft-ietf-avt-ulp-23.txt, Aug. 2007). The payloads of the protected source packets are split into consecutive byte ranges, starting from the beginning of the payload. The first byte range, starting from the beginning of the packet, corresponds to the strongest level of protection, and the protection level decreases as a function of byte range order. Hence, the media data in the protected packets can be organized in such a way that the data appears in descending order of importance within the payload and a similar number of bytes correspond to similar subjective impact in quality among the protected packets. The number of protected levels in FEC repair packets is selectable, and an uneven level of protection is obtained when the number of levels protecting a set of source packets is varied. For example, if there are three levels of protection, one FEC packet may protect all three levels, a second packet may protect the two first levels, and a third packet may protect

only the first level. When applied to RTP payloads containing AMR-WB coded data, the ULP Internet Draft can be used to protect class A bits more robustly compared to class B bits. Details concerning unequal error protection and bit classification of AMR and AMR-WB frames can be found in section 3.6 of IETF RFC 4867, which can be found at www.ietf.org/rfc/rfc4867.txt, and the Third Generation Partnership Project (3GPP) specification 3GPP TS 26.201.

[0019] Another emerging trend in the field of media coding involves what are referred to as “layered codecs.” Such layered codecs include, for example, the International Telecommunications Union (ITU) Telecommunication Standardization Sector (ITU-T), Embedded Variable Bit-Rate (EV-VBR) speech/audio codec and the ITU-T Scalable Video Codec (SVC). The scalable media data comprises a core layer, which is always needed to enable reconstruction at the receiving end, and one or more enhancement layers that can be used to provide added value to the reconstructed media (e.g., to provide improved media quality or an increased level of robustness against transmission errors, etc). The scalability can be deployed at the transmission level, e.g., for controlling the network capacity or shaping a multicast media stream in order to facilitate operation with participants behind access links of different bandwidth. At the application level, the scalability can be utilized for controlling, e.g., computational complexity, encoding delay, or desired quality level. It should be noted that, while in some scenarios the scalability can be applied at the transmitting end-point, there are also operating scenarios where it makes more sense to permit an intermediate network element to perform the scaling.

[0020] From a media transport point of view, scalable codecs provide two basic options--the enhancement layers can be transmitted either together in the same packets with the core layer data, or the (subsets of) enhancement layers can be transmitted in separate packet stream(s). The transmission in separate packet streams also requires a signalling mechanism that can be used to bind together the packet data streams carrying layers of the same media source.

[0021] The approach of carrying all layers (of a media frame) in a single packet provides for a low overhead and easy cross-layer synchronization. On the other hand,

this approach results in more complex scaling in an intermediate network element; the scaling possibility requires the network element to be aware of the details of the packet structure. Furthermore, the scaling also implies modification of packets.

[0022] The approach employing separate data streams for (subsets of) layers provides for a simple scaling possibility because the scaling can be realized by discarding all packets of some data streams. This does not require in-depth knowledge about the packet structure, as long as the (signalling) information on the relationships between the data streams is available. However, this approach results in increased overhead, since each data stream introduces its own protocol overhead (e.g. IP/UDP/RTP). A further challenge associated with this approach involves cross-layer synchronization, i.e., how the receiver can reconstruct the media frames based on layers it receives distributed across multiple data streams. It should be noted that employing multiple data streams, e.g., multiple RTP sessions, is the traditional method for transmitting layered media data within the RTP framework (often referred to as scalable multicast).

[0023] The timeline of a media stream received in RTP packets can be reconstructed based time stamp (TS) information included in the RTP header. The RTP TS provides information on the temporal difference compared to the other RTP packets transmitted in the same RTP session, which enables the placement of each received media frame in its correct place in the timeline. However, the initial value of the RTP TS of an RTP session is random, which implies that RTP TS does not indicate an absolute time. Instead, the RTP TS only references a timing reference within the RTP session. It should be noted that this "randomness" can be considered to be an unknown offset from the absolute time, which is different in each RTP session. Therefore two or more RTP sessions cannot be synchronized based solely on their RTP TS values. This is also true for separate RTP sessions used to carry (subsets of) layers of layered encoding.

[0024] A conventional mechanism within RTP for synchronizing multiple RTP sessions is based on RTCP reports transmitted within each session. In this approach, the sender includes both the timing reference (NTP) and the sending instant in the

RTP TS domain in the RTCP Sender Reports (SR) that it transmits according to specified rules. This enables the receiver to compute the RTP TS offset from the timing reference (NTP) for each of the RTP sessions it receives. These offset values can then be used to match the timing of the media data received in separate RTP sessions, for example to combine layers of a media frame received in multiple RTP sessions. However, this approach requires the first RTCP SRs for each of the RTP sessions to be available before the full reconstruction of a media frame is possible. In practice, this implies that only the core layer of the layered encoding is available until the synchronization information is available.

[0025] One alternative to the above involves pre-synchronizing the RTP TS spaces across RTP sessions in the transmitting end-point. In this approach, the "random" initial value of the RTP TS would be the same for each of the RTP sessions. While this may provide a simple cross-session synchronization mechanism without having to transmit additional data, it is not fully in line with the RTP specification, and existing RTP implementations may not support it. Furthermore, such a mechanism would provide the (pre-)synchronization at the RTP (header) level, but it would be available only for subset of RTP payloads. Such payload type-dependent processing at RTP level is not a desirable feature in a system handling multiple payload types.

[0026] Still another approach for synchronizing multiple RTP sessions involves attaching additional information for each transmitted layer, indicating its temporal location in the presentation timeline. Such information may comprise, for example, a cross-layer sequence number or additional timestamp information that can be used to reconstruct the presentation order of media frames and layers within the frames. However, this approach would still introduce additional overhead for each transmitted layer. This is important because, particularly in the case of smaller pieces of speech/audio data, e.g., in the order of 10-20 bytes, even one additional byte of overhead may have a significant effect on the overall system performance.

SUMMARY OF THE INVENTION

[0027] Various embodiments of the present invention provide systems and methods for providing improved scalable error detection for packet-switched transmission. In one embodiment, common CRC error detection is applied for the core layer and for enhancement layers of the scalable payload in such a way that dropping one or several enhancement layers from the payload does not change the value of the CRC checksum. Only one CRC checksum is transmitted, e.g., in the payload header or in the header of the lower-layer protocol. The transmitter modifies the encoded bit stream in such a manner that the entity in the network deploying the scalable payload and removing layers from the packet does not need to recalculate the CRC placed in the payload or packet header, even when the payload size is changed. A prefix/tail bit field is added in the beginning/end of each enhancement layer to make the CRC check match with the common CRC checksum.

[0028] In another embodiment, the receiver may check the correctness of each received layer simultaneously and, if desired, utilize data only from correctly received layers. Furthermore, the receiver does not need to know the original packet structure, i.e., the number of enhancement layers. The common CRC checksum has the same value for each layer, regardless of the number of transmitted and received enhancement layers.

[0029] Various embodiments also provide for an improved system and method for scalable cross-session timing synchronization. This system and method involves the application of a common CRC checksum for the core layer and the enhancement layers of the scalable media frame in a scenario where the core layer and enhancement layers are transmitted in separate RTP sessions. The transmitted data is formulated in such a way that the CRC checksum computed over an enhancement layer corresponding to a certain core layer is modified to have the same value as the CRC checksum included in the packet carrying the data for the corresponding core layer. The receiver can associate an enhancement layer to the particular core layer by computing the CRC checksum over the enhancement layer and finding a received core layer with the matching CRC checksum. Once a matching pair of core and

enhancement layers is found, the synchronization between RTP sessions carrying the core layer and this particular enhancement layer is established.

[0030] Various embodiments of the present invention can be applied to emerging scalable speech, audio coders and video coders under ITU-T, for example, EV-VBR and SVC.

[0031] These and other advantages and features of the invention, together with the organization and manner of operation thereof, will become apparent from the following detailed description when taken in conjunction with the accompanying drawings, wherein like elements have like numerals throughout the several drawings described below.

BRIEF DESCRIPTION OF THE DRAWINGS

[0032] Figure 1 is an overview diagram of a system within which various embodiments of the present invention may be implemented;

[0033] Figures 2(a) and 2(b) show payload structures with common CRC bit field (polynomial $R_{ref}(z)$) in the payload header and prefix bits (polynomial $R_i(z) + R_{ref}(z)$) in the beginning of each scalable layer;

[0034] Figure 3 is a flow chart showing a process by which various embodiments of the present invention may be implemented;

[0035] Figure 4 is a representation of a payload structure with a common CRC bitfield (polynomial $R_{ref}(z)$) in the core layer packet payload header and tail bits (polynomial $R_i(z) + R_{ref}(z)$) in the end of each scalable layer;

[0036] Figure 5 is a representation of a payload structure when two frames are encapsulated in the same RTP packet;

[0037] Figure 6 is a flow chart showing a process by which other various embodiments of the present invention may be implemented;

[0038] Figure 7 is a perspective view of an electronic device that can be used in conjunction with the implementation of various embodiments of the present invention; and

[0039] Figure 8 is a schematic representation of the circuitry which may be included in the electronic device of Figure 7.

DETAILED DESCRIPTION OF VARIOUS EMBODIMENTS

[0040] Figure 1 is a graphical representation of a generic multimedia communication system within which various embodiments of the present invention may be implemented. As shown in Figure 1, a data source 100 provides a source signal in an analog, uncompressed digital, or compressed digital format, or any combination of these formats. An encoder 110 encodes the source signal into a coded media bitstream. It should be noted that a bitstream to be decoded can be received directly or indirectly from a remote device located within virtually any type of network. Additionally, the bitstream can be received from local hardware or software. The encoder 110 may be capable of encoding more than one media type, or more than one encoder 110 may be required to code different media types of the source signal. The encoder 110 may also get synthetically produced input, such as graphics and text, or it may be capable of producing coded bitstreams of synthetic media. In the following, only processing of one coded media bitstream of one media type is considered to simplify the description. It should be noted, however, that typically real-time broadcast services comprise several streams (typically at least one audio, video and text sub-titling stream). It should also be noted that the system may include many encoders, but in Figure 1 only one encoder 110 is represented to simplify the description without a lack of generality. It should be further understood that, although text and examples contained herein may specifically describe an encoding process, one skilled in the art would understand that the same concepts and principles also apply to the corresponding decoding process and vice versa.

[0041] The coded media bitstream is transferred to a storage 120. The storage 120 may comprise any type of mass memory to store the coded media bitstream. The format of the coded media bitstream in the storage 120 may be an elementary self-contained bitstream format, or one or more coded media bitstreams may be encapsulated into a container file. Some systems operate “live”, i.e. omit storage and

transfer coded media bitstream from the encoder 110 directly to the sender 130. The coded media bitstream is then transferred to the sender 130, also referred to as the server, on a need basis. The format used in the transmission may be an elementary self-contained bitstream format, a packet stream format, or one or more coded media bitstreams may be encapsulated into a container file. The encoder 110, the storage 120, and the sender 130 may reside in the same physical device or they may be included in separate devices. The encoder 110 and sender 130 may operate with live real-time content, in which case the coded media bitstream is typically not stored permanently, but rather buffered for small periods of time in the content encoder 110 and/or in the sender 130 to smooth out variations in processing delay, transfer delay, and coded media bitrate.

[0042] The sender 130 sends the coded media bitstream using a communication protocol stack. The stack may include, but is not limited to, RTP, UDP and IP, although it is also noted that 3GPP circuit-switched telephony may also be used in the context of various embodiments of the present invention. When the communication protocol stack is packet-oriented, the sender 130 encapsulates the coded media bitstream into packets. For example, when RTP is used, the sender 130 encapsulates the coded media bitstream into RTP packets according to an RTP payload format. Typically, each media type has a dedicated RTP payload format. It should be again noted that a system may contain more than one sender 130, but for the sake of simplicity, the following description only considers one sender 130.

[0043] The sender 130 may or may not be connected to a gateway 140 through a communication network. The gateway 140 may perform different types of functions, such as translation of a packet stream according to one communication protocol stack to another communication protocol stack, merging and forking of data streams, and manipulation of data streams according to the downlink and/or receiver capabilities, such as controlling the bit rate of the forwarded stream according to prevailing downlink network conditions. Examples of gateways 140 include MCUs, gateways between circuit-switched and packet-switched video telephony, Push-to-talk over Cellular (PoC) servers, IP encapsulators in digital video broadcasting-handheld

(DVB-H) systems, or set-top boxes that forward broadcast transmissions locally to home wireless networks. When RTP is used, the gateway 140 may be called an RTP mixer or an RTP translator, i.e., an “intermediate node” of an end-to-end RTP session. Alternatively, the gateway 140 may act as an endpoint of an RTP session.

[0044] The system includes one or more receivers 150, typically capable of receiving, de-modulating, and de-capsulating the transmitted signal into a coded media bitstream. The coded media bitstream is transferred to a recording storage 155. The recording storage 155 may comprise any type of mass memory to store the coded media bitstream. The recording storage 155 may alternatively or additively comprise computation memory, such as random access memory. The format of the coded media bitstream in the recording storage 155 may be an elementary self-contained bitstream format, or one or more coded media bitstreams may be encapsulated into a container file. If there are many coded media bitstreams associated with each other, a container file is typically used and the receiver 150 comprises or is attached to a container file generator producing a container file from input streams. Some systems operate “live,” i.e., omit the recording storage 155 and transfer coded media bitstream from the receiver 150 directly to the decoder 160. In some systems, only the most recent part of the recorded stream, e.g., the most recent 10-minute excerpt of the recorded stream, is maintained in the recording storage 155, while any earlier recorded data is discarded from the recording storage 155.

[0045] The coded media bitstream is transferred from the recording storage 155 to the decoder 160. If there are many coded media bitstreams associated with each other and encapsulated into a container file, a file parser (not shown in the figure) is used to decapsulate each coded media bitstream from the container file. The recording storage 155 or a decoder 160 may comprise the file parser, or the file parser is attached to either recording storage 155 or the decoder 160.

[0046] The coded media bitstream is typically processed further by a decoder 160, whose output is one or more uncompressed media streams. Finally, a renderer 170 may reproduce the uncompressed media streams with a loudspeaker, for example.

The receiver 150, recording storage 155, decoder 160, and renderer 170 may reside in the same physical device or they may be included in separate devices.

[0047] Various embodiments of the present invention provide systems and methods for providing improved scalable error detection for packet-switched transmission. In one embodiment, common CRC error detection is applied for the core layer and for enhancement layers of the scalable payload in such a way that dropping one or several enhancement layers from the payload does not change the value of the CRC checksum. Only one CRC checksum is transmitted, e.g., in the payload header or in the header of the lower-layer protocol. The transmitter modifies the encoded bit stream in such a manner that the entity in the network deploying the scalable payload and removing layers from the packet does not need to recalculate the CRC placed in the payload or packet header, even when the payload size is changed. A prefix/tail bit field is added in the beginning/end of each enhancement layer to make the CRC check match with the common CRC checksum.

[0048] In another embodiment, the receiver may check the correctness of each received layer simultaneously and, if desired, utilize data only from correctly received layers. Furthermore, the receiver does not need to know the original packet structure, i.e., the number of transmitted enhancement layers. The common CRC checksum has the same value for each layer, regardless of the number of transmitted and received enhancement layers.

[0049] In various embodiments of the present invention, having a common CRC check polynomial for core and each enhancement layer, a network element or another entity may remove enhancement layer(s) from the payload without touching the CRC and hence without a need to recalculate the CRC covering the payload. Furthermore, the receiver does not need to know whether the payload was touched between the transmitter and receiver. The CRC verification will match, regardless of how many enhancement layers are received. The common CRC is especially advantageous for scalable content, when the CRC is located in the (lower-layer) packet header, and the CRC check is conducted over the whole payload.

[0050] According to various embodiments, having a separate CRC check for each enhancement layer is not necessary. In such a case, no separate error check for each layer is conducted. However, when separate error checks can be conducted, the receiver may reconstruct the frame using only the correctly received core layer and enhancement layers. Furthermore, unequal error protection can be used for the core layer and enhancement layers. For example, when the core bitstream is received error free, the remaining part of the bit stream having less or no protection at all can be applied to enhance the decoding and to reconstruct the higher bit rate mode. When the common CRC is placed in the lower-layer protocol header, a network element may remove some of the enhancement layers from the payload without a need for recalculating the checksum information in the header. This can be particularly useful because, while some applications/codecs can use enhancement layer data that may or may not contain errors, others are not capable of doing so.

[0051] In more conventional systems, the media payload header contains the checksum calculated over the whole payload (equal error detection) for regular UDP transport, or partially over the most sensitive part of the payload (unequal error detection) for e.g., UDP-Lite transport. Furthermore, the Datagram Congestion Control Protocol (DCCP), which is discussed at www.ietf.org/rfc/rfc4340.txt, could be used to carry the CRC checksum as an (optional) part of the protocol header.

[0052] The CRC checksum is the most common error detection method. The checksum is generated by basically dividing the message polynomial with a selected and well known key polynomial. The remainder of the polynomial division operation is the actual CRC checksum information that will be transmitted.

$$M(z) = Q(z)K(z) + R(z)$$

[0053] In the above, $M(z)$ is the actual message in the payload, $K(z)$ is the key polynomial (divisor), $Q(z)$ is the quotient and $R(z)$ is the remainder, which is also the CRC checksum transmitted to the receiver as part of the packet (either in the payload or in the packet header). When the length of the key polynomial $K(z)$ is $N + 1$, then the length of the remainder (CRC checksum) is N . The following shows a simple CRC calculation:

$$M(z)=[1\ 1\ 0\ 1\ 1\ 0\ 1]\ (=z^6+z^5+z^3+z^2+z^0),\ K(z)=[1\ 0\ 1\ 1]\ (=z^3+z^1+z^0)$$

$$\begin{array}{cccccccc}
 & & & & 1 & 1 & 1 & 1 \\
 1 & 0 & 1 & 1 & \overline{) 1} & 1 & 0 & 1 & 1 & 0 & 1 \\
 & & & & 1 & 0 & 1 & 1 & & & \\
 & & & & \hline
 & & & & 1 & 1 & 0 & 1 & & & \\
 & & & & 1 & 0 & 1 & 1 & & & \\
 & & & & \hline
 & & & & 1 & 1 & 0 & 0 & & & \\
 & & & & 1 & 0 & 1 & 1 & & & \\
 & & & & \hline
 & & & & 1 & 1 & 1 & 1 & & & \\
 & & & & 1 & 0 & 1 & 1 & & & \\
 & & & & \hline
 & & & & & 1 & 0 & 0 & & &
 \end{array}$$

[0054] In the above, the remainder (CRC polynomial) of the calculation is $R(z) = [1 \ 0 \ 0]$.

[0055] Alternatively, the CRC check can be conducted by appending the message polynomial with zeros prior to CRC computation. In other words, a zero vector of the length of the CRC polynomial is appended to the end of the message. The message bits are effectively shifted left by N places, and N zeros are inserted in the least significant bits of the message, i.e., $M(z)$ is multiplied by z^N , and the outcome is as follows:

$$M(z)z^N = Q(z)K(z) + R(z)$$

[0056] The following shows a simple CRC calculation illustrating the effect of appending a zero vector to the end of the message:

$$M(z)=[1\ 1\ 0\ 1\ 1\ 0\ 1]\ (=z^6+z^5+z^3+z^2+z^0),\ K(z)=[1\ 0\ 1\ 1]\ (=z^3+z^1+z^0).$$

[0057] The message polynomial $M(z)$ is then multiplied with z^3 to determine the remainder. Hence, $M(z) z^3 = [1 \ 1 \ 0 \ 1 \ 1 \ 0 \ 1 \ 0 \ 0 \ 0]$

[illegible]

$$\begin{array}{r}
 1 \ 0 \ 0 \ 0 \\
 1 \ 0 \ 1 \ 1 \\
 \hline
 0 \ 1 \ 1 \ 0 \\
 0 \ 0 \ 0 \ 0 \\
 \hline
 1 \ 1 \ 0 \ 0 \\
 1 \ 0 \ 1 \ 1 \\
 \hline
 1 \ 1 \ 1
 \end{array}$$

[0058] In the above, the remainder (CRC polynomial) of the calculation is $R(z) = [1 \ 1 \ 1]$.

[0059] If the remainder polynomial is appended to both sides of the equation $M(z)z^N = Q(z)K(z) + R(z)$, the outcome is that the remainder polynomial of length N is appended at the end of the message polynomial, and the remainder of the division operation by the key polynomial disappears:

$$M(z)z^N + R(z) = Q(z)K(z)$$

[0060] An example showing the remainder $R(z)$ appearing as zero when the CRC checksum polynomial is appended at the end of the message polynomial is as follows:

$$M(z) = [1 \ 1 \ 0 \ 1 \ 1 \ 0 \ 1], K(z) = [1 \ 0 \ 1 \ 1], R(z) = [1 \ 1 \ 1]$$

$$\begin{array}{r}
 1 \ 1 \ 1 \ 1 \ 1 \ 0 \ 1 \\
 1 \ 0 \ 1 \ 1 \ \bigg| \begin{array}{r}
 1 \ 1 \ 0 \ 1 \ 1 \ 0 \ 1 \ 1 \ 1 \ 1 \\
 1 \ 0 \ 1 \ 1 \\
 \hline
 1 \ 1 \ 0 \ 1 \\
 1 \ 0 \ 1 \ 1 \\
 \hline
 1 \ 1 \ 0 \ 0 \\
 1 \ 0 \ 1 \ 1 \\
 \hline
 1 \ 1 \ 1 \ 1 \\
 1 \ 0 \ 1 \ 1 \\
 \hline
 1 \ 0 \ 0 \ 1 \\
 1 \ 0 \ 1 \ 1 \\
 \hline
 0 \ 1 \ 0 \ 1 \\
 0 \ 0 \ 0 \ 0 \\
 \hline
 1 \ 0 \ 1 \ 1 \\
 1 \ 0 \ 1 \ 1 \\
 \hline
 0 \ 0 \ 0
 \end{array}
 \end{array}$$

[0061] This result is applied later when designing a common CRC polynomial for a scalable bit stream. The advantage of the result is that the receiver, in order to

conduct the CRC check, does not need to compare the outcome of the polynomial division with the transmitted CRC bits, but the error check is completed and error free transmission is verified or determined when the remainder is zero.

[0062] According to various embodiments, a scalable bitstream is accompanied by a common CRC field. The payload header or the lower-level protocol header contains the CRC field, against which the correctness of the received payload is checked. The same CRC field is utilized for the core layer as well as to all enhancement layers.

[0063] In general, the CRC check for the core layer bit stream $M_{core}(z)$ is as follows:

$$M_{core}(z) = Q_{core}(z)K(z) + R_{ref}(z)$$

[0064] The remainder polynomial $R_{ref}(z)$ computed for the core layer is used as the CRC check for the enhancement layers as well. When dividing the message polynomial containing both the core and the enhancement layer(s), the remainder is different. It should be noted that the core layer and the enhancement layers can be handled together as a single bitstream, or they can be handled separately. Therefore, prefix/tail bits need to be inserted in the beginning/end of each enhancement layer.

[0065] When the remainder (CRC check bits) computed for the core layer bit stream is $R_{ref}(z)$, any arbitrary bit stream can be tailored to produce the same remainder.

The method involves adding the remainder polynomial of the core layer bit stream as the prefix/tail bits to the message polynomial, i.e., (bit-wise) adding $R_{ref}(z)$ to the $R(z)$ appended to the end of the message polynomial. The outcome is the desired remainder as follows:

$$M(z)z^N + R(z) + R_{ref}(z) = Q(z)K(z) + R_{ref}(z)$$

[0066] Instead of having a zero remainder, the CRC check of the message $M(z)z^N + R(z) + R_{ref}(z)$ will produce the desired output remainder polynomial $R_{ref}(z)$. It should also be noted that the message length is not increased when the desired remainder polynomial is added at the end of the message since the length of the prefix/tail bit field $R(z) + R_{ref}(z)$ is the same N .

[0067] The following example shows the desired remainder appearing as the outcome of the division operation when it is first added to the CRC polynomial at the end of the message:

$M(z) = [1\ 1\ 0\ 1\ 1\ 0\ 1]$, $K(z) = [1\ 0\ 1\ 1]$, $R(z) = [1\ 1\ 1]$, desired remainder for example $R_{ref}(z) = [1\ 0\ 0]$, $R(z) + R_{ref}(z) = [0\ 1\ 1]$

$$\begin{array}{r}
 \\
 1\ 0\ 1\ 1\ 1 \overline{) \begin{array}{l} 1\ 1\ 0\ 1\ 1\ 0\ 1\ 0\ 1\ 1 \\ 1\ 0\ 1\ 1 \\ \hline 1\ 1\ 0\ 1 \\ 1\ 0\ 1\ 1 \\ \hline 1\ 1\ 0\ 0 \\ 1\ 0\ 1\ 1 \\ \hline 1\ 1\ 0\ 1 \\ 1\ 1\ 1\ 1 \\ \hline 1\ 0\ 1\ 1 \\ 1\ 0\ 0\ 0 \\ \hline 1\ 0\ 1\ 1 \\ 0\ 1\ 1\ 1 \\ \hline 0\ 0\ 0\ 0 \\ 1\ 1\ 1\ 1 \\ \hline 1\ 0\ 1\ 1 \\ 1\ 0\ 0 \end{array} }
 \end{array}$$

[0068] As indicated above, the desired remainder $R_{ref}(z) = [1\ 0\ 0]$ appears in the calculation.

[0069] It is helpful to consider a case where the payload comprises a core bit stream and m enhancement layers. The length of the core is N_0 bits and the length of the enhancement layer i is N_i bits. Thus, the message polynomial is as follows:

$$M_{payload}(z) = M_m(z)z^{N_0+N_1+\dots+N_{m-1}} + M_{m-1}(z)z^{N_0+N_1+\dots+N_{m-2}} + \dots + M_1(z)z^{N_0} + M_{core}(z)$$

[0070] Instead of the dedicated remainder $R_i(z)$ of each enhancement layer, the prefix/tail bits comprising the polynomial $R_i(z) + R_{ref}(z)$ are added in the beginning/end of each enhancement layer $M_i(z)$. The result is that a common CRC polynomial $R_{ref}(z)$ can be utilized for checking the correctness of the core and each enhancement layer separately. A CRC check (in the receiver) over the core layer and

any of the enhancement layers produces an identical result. In other words, the remainder of the division operation can always be compared against the CRC bit field, which can be located in the payload, in the payload header or in the packet header in various embodiments.

[0071] Figures 2(a) and 2(b) show a payload structure with common CRC bit field 210 (polynomial $R_{ref}(z)$) in a payload header 200 and prefix bits 240 (polynomial $R_i(z) + R_{ref}(z)$) in the beginning of each scalable layer 220. A core layer 230 is also shown. A CRC check of each enhancement layer 220, either separately (as shown in Figure 2(a)) or together with all lower layers (as shown in Figure 2(b)), will produce the common remainder stored in the payload (or packet) header 200. Additionally, the CRC field 210 could be carried somewhere else in the payload, in the packet header or the header of a lower-level protocol instead of in the payload header. Still further, the CRC check can be conducted as a combination or hybrid of the systems depicted in Figures 2(a) and 2(b). For example, up to a certain point, the CRC check can be performed for each respective layer together with all lower layers. After this point, however, a separate check can be performed on each layer individually (without considering lower layers). Other variations of such a hybrid system are also possible.

[0072] Figure 3 is a flow chart showing a process by which various embodiments of the present invention may be implemented. At 310 in Figure 3, a plurality of layers are encoded, i.e., formed into a packet, including a core layer and at least one enhancement layer. At 320, a single checksum is encoded. It should be noted that, although 310 and 320 are shown as occurring sequentially, these processes can occur simultaneously and/or as a single process. As discussed previously, the single checksum is applicable for use in performing error checks of the core layer and each enhancement layer. In the case of the enhancement layers, the single checksum is applicable in error checks for each enhancement layer either individually or in combination with each lower enhancement layer and the core layer. As discussed previously, this may be accomplished when necessary by adding an additional field to each enhancement layer such that the single checksum will be matched during any subsequent error check where no errors or missing data are present. The adding of

these fields is represented at 300. At 330, the encoded data is transmitted from a sending apparatus to a receiving apparatus.

[0073] At 340, the core layer and any enhancement layers are processed from the bitstream by the receiver, and the single checksum is processed at 350. Again, although these processes are shown as occurring sequentially, these processes can occur simultaneously and/or as a single process. At 360, at least one error check is performed for one or more of the layers using the single checksum. Regardless of which layer(s) are being checked, the single checksum is usable for the error check. As discussed previously, it is possible for each enhancement layer to be separately checked for errors, or each enhancement layer can be checked in combination with each lower enhancement layer and the core layer. At 370 and in one embodiment, only those layers where no data was missing or in error is used for the further processing.

[0074] Various embodiments also provide for an improved system and method for scalable cross-session timing synchronization. This system and method involves the application of a common CRC checksum for the core layer and the enhancement layers of the scalable media frame in a scenario where the core layer and enhancement layers are transmitted in separate RTP sessions. The transmitted data is formulated in such a way that the CRC checksum computed over an enhancement layer corresponding to a certain core layer is modified to have the same value as the CRC checksum included in the packet carrying the data for the corresponding core layer. The receiver can associate an enhancement layer to the particular core layer by computing the CRC checksum over the enhancement layer and finding a received core layer with the matching CRC checksum. Once a matching pair of core and enhancement layers is found, the synchronization between RTP sessions carrying the core layer and this particular enhancement layer is established.

[0075] Figure 4 presents a typical transmission system for scalable media content, with a common CRC bit field 210 (polynomial $R_{ref}(z)$) in the core layer packet payload header 200 and tail bits 400 (polynomial $R_i(z) + R_{ref}(z)$) in the end of each scalable layer 220. The core layer 230 and each enhancement layer 220 are conveyed

in separate RTP packets in separate RTP sessions. The encoder/transmitter determines a tail bit field 400 for each enhancement layer 220. The tail bits 400, as well as the corresponding core layer CRC checksum 210, are attached to either an RTP payload, a payload header or a packet header. Alternatively, the CRC checksum 210 can be attached to any of the enhancement layer packets. A CRC check of each layer separately will produce an identical checksum with the one transmitted in the payload, payload header or packet header of the core layer.

[0076] Figure 5 presents an alternative transmission mechanism for scalable media content transmitting more than one frame in an RTP packet. In this arrangement, the core layer 230 and each enhancement layer 220 are conveyed in separate RTP packets, i.e., in separate RTP sessions. A common CRC bit field 210 (polynomial $R_{ref}(z)$) is placed in the core layer packet payload header, and tail bits 400 (polynomial $R_i(z) + R_{ref}(z)$) are placed in the end of the core layer 230 of the second frame and in each scalable layer. The encoder/transmitter determines a tail bit field 400 as described above. The tail bits 400, as well as the corresponding core layer CRC checksum 210, are attached to an RTP payload, a payload header or a packet header. Alternatively, the CRC checksum 210 can be attached to any of the enhancement layer packets. Once again, a CRC check of each layer separately will produce an identical checksum with the one transmitted in the payload, payload header or packet header of the core layer 230.

[0077] It should be noted that the CRC checksum of the second frame of a payload/packet is not critical for synchronization purposes in various embodiments. When the synchronization of the core layer and enhancement layers of the first frame of the payload/packet is achieved, the second frame is automatically in synchronization. Instead, the CRC checksum can be utilized for scalable unequal error detection only, as described previously.

[0078] In an alternative to the above, the tail bits can be attached in the encoder/transmitter in such a manner that the checksum is always combined with the previous enhancement layer and core layer. In other words, the checksum of the core layer + enhancement layer #1 needs to match with the core layer CRC. Furthermore,

the checksum of the core layer + enhancement layer #1 + enhancement layer #2 + ... + enhancement layer #N needs to match with the core layer CRC. This embodiment prevents the usage of higher layer enhancement layers when an intermediate enhancement layer is missing. This can be particularly advantageous in the event of an incremental architecture in which additional enhancement layers always utilize the information of all lower layers.

[0079] The receiver synchronizes the core layer with the corresponding enhancement layers by comparing the CRC checksum computed at the receiver at each layer (received in separate RTP sessions) with the received CRC checksum for the core layer. If the checksum does not match, then the enhancement layer RTP session is either corrupted or it is supposed to be attached to another core layer RTP session. If the checksum does match, then the enhancement layer belongs to the same media frame with the core layer, and the RTP TS offset between the RTP session for the core layer and the RTP session for the enhancement layer at issue can be computed. This provides the necessary information for cross-layer synchronization between these layers. It should be noted that re-synchronization is needed only in the event that a discontinuity is introduced for the timestamp space in either of the RTP sessions. Thus, it is sufficient to perform the cross-layer synchronization, for example, in a few frames in the beginning of the session or (in a few frames) at regular intervals.

[0080] It should also be noted that when the common CRC is placed in the lower-layer protocol header, a network element may remove some of the enhancement layers from the payload without having to recalculate the checksum information in the header.

[0081] By having a common CRC check polynomial for the core layer and each enhancement layer, the synchronization between the core layer and all of the enhancement layers is maintained without having to transmit additional synchronization data. The dedicated bit patterns of the CRC checksum in the core layer (or one of the enhancement layers) and the required tail bits ensure that wrong combinations of core and enhancement layers are not created. Additionally, this

arrangement does not suffer from any IP/UDP protocol limitations, as the synchronization is done in the application layer. Still further, it is possible to reuse the existing CRC checksum based error detection mechanism for the scalable media formats, particularly when more than one frame is encapsulated in the same RTP packet.

[0082] Figure 6 is a flow chart showing a process by which various embodiments of the present invention may be implemented. At 610 in Figure 6, a plurality of layers are encoded, including a core layer and at least one enhancement layer. At 620, a single checksum is encoded. It should be noted that, although 610 and 620 are shown as occurring sequentially, these processes can occur simultaneously and/or as a single process. As discussed previously, the single checksum is applicable for use in performing error checks of the core layer and each corresponding enhancement layer. As discussed previously, this may be accomplished when necessary by adding an additional field (the tail bits fields of Figures 4 and 5) to one or more of the enhancement layers such that the single checksum will be matched during any subsequent error check where no errors or missing data are present. The adding of these fields is represented at 600. The transmitted data is formulated in such a way that the CRC checksum computed over an enhancement layer corresponding to a certain core layer is modified to have the same value as the CRC checksum included in the packet carrying the data for the corresponding core layer. At 630, the encoded data is transmitted from a sending apparatus to a receiving apparatus in multiple sessions.

[0083] At 640, the core layer and enhancement layers are processed from the bitstream by the receiver, and the single checksum is processed at 650. Again, although these processes are shown as occurring sequentially, these processes can occur simultaneously and/or as a single process. At 660, the CRC checksum is determined over an enhancement layer. At 670, the receiver finds a received core layer with a CRC checksum which matches the CRC checksum determined at 660. When such a matching CRC checksum is found, the enhancement layer and the core layer are associated with each other at 680. At this point, the synchronization

between RTP sessions carrying the core layer identified at 670 and the enhancement layer for which the CRC checksum was determined at 660 is established.

[0084] Figures 7 and 8 show one representative mobile device 12 within which the present invention may be implemented. It should be understood, however, that the present invention is not intended to be limited to one particular type of electronic device. The mobile device 12 of Figures 7 and 8 includes a housing 30, a display 32 in the form of a liquid crystal display, a keypad 34, a microphone 36, an ear-piece 38, a battery 40, an infrared port 42, an antenna 44, a smart card 46 in the form of a UICC according to one embodiment of the invention, a card reader 48, radio interface circuitry 52, codec circuitry 54, a controller 56 and a memory 58. Individual circuits and elements are all of a type well known in the art, for example in the Nokia range of mobile telephones.

[0085] The various embodiments of the present invention described herein is described in the general context of method steps or processes, which may be implemented in one embodiment by a computer program product, embodied in a computer-readable medium, including computer-executable instructions, such as program code, executed by computers in networked environments. A computer-readable medium may include removable and non-removable storage devices including, but not limited to, Read Only Memory (ROM), Random Access Memory (RAM), compact discs (CDs), digital versatile discs (DVD), etc. Generally, program modules may include routines, programs, objects, components, data structures, etc. that perform particular tasks or implement particular abstract data types. Computer-executable instructions, associated data structures, and program modules represent examples of program code for executing steps of the methods disclosed herein. The particular sequence of such executable instructions or associated data structures represent examples of corresponding acts for implementing the functions described in such steps or processes.

[0086] Software and web implementations of various embodiments of the present invention can be accomplished with standard programming techniques with rule-based logic and other logic to accomplish various database searching steps or

processes, correlation steps or processes, comparison steps or processes and decision steps or processes. It should be noted that the words “component” and “module,” as used herein and in the following claims, is intended to encompass implementations using one or more lines of software code, and/or hardware implementations, and/or equipment for receiving manual inputs.

[0087] The foregoing description of embodiments of the present invention have been presented for purposes of illustration and description. The foregoing description is not intended to be exhaustive or to limit embodiments of the present invention to the precise form disclosed, and modifications and variations are possible in light of the above teachings or may be acquired from practice of various embodiments of the present invention. The embodiments discussed herein were chosen and described in order to explain the principles and the nature of various embodiments of the present invention and its practical application to enable one skilled in the art to utilize the present invention in various embodiments and with various modifications as are suited to the particular use contemplated. The features of the embodiments described herein may be combined in all possible combinations of methods, apparatus, modules, systems, and computer program products.

WHAT IS CLAIMED IS:

- 1 1. A method for encoding data into a bitstream, comprising:
2 encoding a core layer and at least one corresponding enhancement
3 layer into the bitstream; and
4 encoding a single checksum into the bitstream, the single checksum
5 being applicable for use in performing a check of each corresponding enhancement
6 layer, either alone or in combination with each lower corresponding enhancement
7 layer and the core layer.
- 1 2. The method of claim 1, wherein the single checksum comprises a
2 cyclic redundancy check checksum.
- 1 3. The method of claim 1, wherein the single checksum is encoded in a
2 payload header for the encoded core layer and the at least one corresponding
3 enhancement layer.
- 1 4. The method of claim 1, wherein the single checksum is encoded in a
2 header of a lower-layer protocol for the encoded core layer and the at least one
3 corresponding enhancement layer.
- 1 5. The method of claim 1, wherein the single checksum is encoded in a
2 payload.
- 1 6. The method of claim 1, wherein the single checksum is encoded in a
2 packet header.
- 1 7. The method of claim 1, further comprising, before encoding the at least
2 one corresponding enhancement layer, adding an additional field to at least one of the
3 at least one corresponding enhancement layer such that any subsequent check
4 involving a corresponding enhancement layer including the additional field will match
5 the single checksum if no data is missing or in error.

1 8. The method of claim 7, wherein the additional field is added at the
2 beginning of each corresponding enhancement layer.

1 9. The method of claim 7, wherein the additional field is added at the end
2 of each corresponding enhancement layer.

1 10. The method of claim 7, wherein the additional field is encoded in a
2 payload.

1 11. The method of claim 7, wherein the additional field is encoded in a
2 packet header.

1 12. The method of claim 7, wherein the additional field is encoded in a
2 payload header.

1 13. The method of claim 1, wherein the core layer and the at least one
2 corresponding enhancement layer are transmitted to a remote device in a plurality of
3 sessions.

1 14. A computer program product, embodied in a computer-readable
2 storage medium, comprising computer code configured to perform the processes of
3 claim 1.

1 15. An apparatus, comprising:
2 an electronic device configured to:
3 encode a core layer and at least one corresponding
4 enhancement layer into the bitstream; and
5 encode a single checksum into the bitstream, the single
6 checksum being applicable for use in performing a check of each corresponding
7 enhancement layer, either alone or in combination with each lower corresponding
8 enhancement layer and the core layer..

1 16. The apparatus of claim 15, wherein the single checksum comprises a
2 cyclic redundancy check checksum.

1 17. The apparatus of claim 15, wherein the single checksum is encoded in
2 a payload header for the encoded core layer and the at least one corresponding
3 enhancement layer.

1 18. The apparatus of claim 15, wherein the single checksum is encoded in
2 a header of a lower-layer protocol for the encoded core layer and the at least one
3 corresponding enhancement layer.

1 19. The apparatus of claim 15, wherein the single checksum is encoded in
2 a payload.

1 20. The apparatus of claim 15, wherein the single checksum is encoded in
2 a packet header.

1 21. The apparatus of claim 15, wherein the electronic device is further
2 configured to, before encoding the at least one corresponding enhancement layer, add
3 an additional field to at least one of the at least one corresponding enhancement layer
4 such that any subsequent check involving a corresponding enhancement layer
5 including the additional field will match the single checksum if no data is missing or
6 in error.

1 22. The apparatus of claim 21, wherein the additional field is added at the
2 beginning of each corresponding enhancement layer.

1 23. The apparatus of claim 21, wherein the additional field is added at the
2 end of each corresponding enhancement layer.

1 24. The apparatus of claim 21, wherein the additional field is encoded in a
2 payload.

1 25. The apparatus of claim 21, wherein the additional field is encoded in a
2 packet header.

1 26. The apparatus of claim 21, wherein the additional field is encoded in a
2 payload header.

1 27. The apparatus of claim 15, wherein the core layer and the at least one
2 corresponding enhancement layer are transmitted to a remote device in a plurality of
3 sessions.

1 28. An apparatus, comprising:
2 means for encoding a core layer and at least one corresponding
3 enhancement layer into the bitstream; and
4 means for encoding a single checksum into the bitstream, the single
5 checksum being applicable for use in performing a check of each corresponding
6 enhancement layer, either alone or in combination with each lower corresponding
7 enhancement layer and the core layer.

1 29. The apparatus of claim 28, further comprising means for, before
2 encoding the at least one corresponding enhancement layer, adding an additional field
3 to at least one of the at least one corresponding enhancement layer such that any
4 subsequent check involving a corresponding enhancement layer including the
5 additional field will match the single checksum if no data is missing or in error..

1 30. A method of processing encoded data from a bitstream;
2 processing a core layer and at least one corresponding enhancement
3 layer from the bitstream;
4 processing a single checksum from the bitstream; and
5 performing a check of one of the at least one corresponding
6 enhancement layer, either alone or in combination with each lower corresponding
7 enhancement layer and the core layer, using the single checksum, wherein the single

8 checksum is usable regardless of the number of corresponding enhancement layers
9 which are used in the check.

1 31. The method of claim 30, wherein the single checksum comprises a
2 cyclic redundancy check checksum.

1 32. The method of claim 30, wherein the single checksum is located in a
2 payload header for the encoded core layer and the at least one corresponding
3 enhancement layer.

1 33. The method of claim 30, wherein the single checksum is located in a
2 header of a lower-layer protocol for the encoded core layer and the at least one
3 corresponding enhancement layer.

1 34. The method of claim 30, wherein the single checksum is encoded in a
2 payload.

1 35. The method of claim 30, wherein the single checksum is encoded in a
2 packet header.

1 36. The method of claim 30, wherein an additional field has been
2 previously added to at least one of the at least one corresponding enhancement layer
3 such that the check involving the respective enhancement layer will match the single
4 checksum if no data is missing or in error.

1 37. The method of claim 36, wherein the additional field has been added at
2 the beginning of each corresponding enhancement layer.

1 38. The method of claim 36, wherein the additional field has been added at
2 the end of each corresponding enhancement layer.

1 39. The method of claim 36, wherein the additional field is encoded in a
2 payload.

1 40. The method of claim 36, wherein the additional field is encoded in a
2 packet header.

1 41. The method of claim 36, wherein the additional field is encoded in a
2 payload header.

1 42. The method of claim 30, wherein the core layer and the at least one
2 corresponding enhancement layer are received from a remote device in a plurality of
3 sessions, and further comprising:

4 using the performed check to match the core layer with the at least one
5 corresponding enhancement layer; and

6 synchronizing the plurality of sessions including the core layer and the
7 at least one corresponding enhancement layer.

1 43. The method of claim 30, wherein checks are performed simultaneously
2 for each corresponding enhancement layer and the core layer, and further comprising
3 using in a subsequent process only those layers which have been correctly received.

1 44. A computer program product, embodied in a computer-readable
2 storage medium, comprising computer code configured to perform the processes of
3 claim 30.

1 45. An apparatus, comprising:

2 an electronic device configured to:

3 process a core layer and at least one corresponding
4 enhancement layer from the bitstream;

5 process a single checksum from the bitstream; and

6 perform a check of one of the at least one corresponding

7 enhancement layer, either alone or in combination with each lower corresponding

8 enhancement layer and the core layer, using the single checksum, wherein the single

9 checksum is usable regardless of the number of corresponding enhancement layers

10 which are used in the check

1 46. The apparatus of claim 45, wherein the single checksum comprises a
2 cyclic redundancy check checksum.

1 47. The apparatus of claim 45, wherein the single checksum is located in a
2 payload header for the encoded core layer and the at least one corresponding
3 enhancement layer.

1 48. The apparatus of claim 45, wherein the single checksum is located in a
2 header of a lower-layer protocol for the encoded core layer and the at least one
3 corresponding enhancement layer.

1 49. The apparatus of claim 45, wherein the single checksum is encoded in
2 a payload.

1 50. The apparatus of claim 45, wherein the single checksum is encoded in
2 a packet header.

1 51. The apparatus of claim 45, wherein an additional field has been
2 previously added to at least one of the at least one corresponding enhancement layer
3 such that the check involving the respective corresponding enhancement layer will
4 match the single checksum if no data is missing or in error.

1 52. The apparatus of claim 51, wherein the additional field has been added
2 at the beginning of each corresponding enhancement layer.

1 53. The apparatus of claim 51, wherein the additional field has been added
2 at the end of each corresponding enhancement layer.

1 54. The apparatus of claim 51, wherein the additional field is encoded in a
2 payload.

1 55. The apparatus of claim 51, wherein the additional field is encoded in a
2 packet header.

1 56. The apparatus of claim 51, wherein the additional field is encoded in a
2 payload header.

1 57. The apparatus of claim 45, wherein the core layer and the at least one
2 corresponding enhancement layer are received from a remote device in a plurality of
3 sessions, and wherein the electronic device is further configured to:
4 use the performed check to match the core layer with the at least one
5 corresponding enhancement layer; and
6 synchronize the plurality of sessions including the core layer and the at
7 least one corresponding enhancement layer.

1 58. The apparatus of claim 45, wherein checks are performed
2 simultaneously for each corresponding enhancement layer and the core layer, and
3 wherein the electronic device is further configured to use in a subsequent process only
4 those layers which have been correctly received.

1 59. An apparatus, comprising:
2 means for processing a core layer and at least one corresponding
3 enhancement layer from the bitstream;
4 means for processing a single checksum from the bitstream; and
5 means for performing a check of one of the at least one corresponding
6 enhancement layer, either alone or in combination with each lower corresponding
7 enhancement layer and the core layer, using the single checksum, wherein the single
8 checksum is usable regardless of the number of corresponding enhancement layers
9 which are used in the check.

1 60. The apparatus of claim 35, wherein an additional field has been
2 previously added to at least one of the at least one corresponding enhancement layer
3 such that the check involving the respective corresponding enhancement layer will
4 match the single checksum if no data is missing or in error.

- 1 61. The apparatus of claim 60, wherein checks are performed
2 simultaneously for each corresponding enhancement layer and the core layer, and
3 further comprising means for using in a subsequent process only those layers which
4 have been correctly received.

1/8




FIG. 1

FIG. 2(a)





FIG. 2(b)



3 / 8




FIG. 3





FIG. 5



6/8




FIG. 6




FIG. 7




FIG. 8