



(51) International Patent Classification:

H04N 19/105 (2014.01) *H04N 19/176* (2014.01)
H04N 19/119 (2014.01) *H04N 19/543* (2014.01)
H04N 19/137 (2014.01)

(21) International Application Number:

PCT/US2016/067792

(22) International Filing Date:

20 December 2016 (20.12.2016)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

15/010,594 29 January 2016 (29.01.2016) US

(71) Applicant: **GOOGLE INC.** [US/US]; 1600 Amphitheatre Parkway, Mountain View, California 94043 (US).

(72) Inventors: **BANKOSKI, James**; c/o Google Inc., 1600 Amphitheatre Parkway, Mountain View, California 94043 (US). **MUKHERJEE, Debargha**; c/o Google Inc., 1600 Amphitheatre Parkway, Mountain View, California 94043 (US). **XU, Yaowu**; c/o Google Inc., 1600 Amphitheatre Parkway, Mountain View, California 94043 (US).


(74) Agents: **BASILE, Jr., Andrew R.** et al.; Young Basile Hanlon & MacFarlane, P.C., 3001 West Big Beaver Road, Suite 624, Troy, Michigan 48084 (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, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, 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, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

[Continued on next page]

(54) Title: MOTION VECTOR PREDICTION USING PRIOR FRAME RESIDUAL



(57) Abstract: A mask is generated from a residual that is a difference between pixel values of at least two frames other than a current frame in a video sequence. The mask is then used to encode or decode a current block by inter-prediction. The mask forms an edge of an object within its bounds and can be used to combine different portions of a reference frame into a single block. This can increase coding efficiency.

FIG. 7



Published:

— *with international search report (Art. 21(3))*

MOTION VECTOR PREDICTION USING PRIOR FRAME RESIDUAL

BACKGROUND

[0001] Digital video streams typically represent video using a sequence of frames or still images. Each frame can include a number of blocks, which in turn may contain information describing the value of color, brightness or other attributes for pixels. The amount of data in a video stream is large, and transmission and storage of video can use significant computing or communications resources. Due to the large amount of data involved in video data, high performance compression is needed for transmission and storage. This often involves inter prediction using motion vectors.

SUMMARY

[0002] This disclosure relates in general to encoding and decoding visual data, such as video stream data, using motion vector prediction using a prior frame residual.

[0003] A method for encoding or decoding a video signal according to one aspect of this disclosure comprises generating a mask for a current block within a current frame from a residual that is a difference between pixel values of at least two frames other than the current frame, and encoding or decoding the current block by inter-prediction using the mask. The video signal includes frames defining a video sequence, the frames having blocks formed of pixels.

[0004] One aspect of an apparatus described herein is an apparatus for encoding or decoding a video signal, the video signal including frames defining a video sequence, the frames having blocks formed of pixels. The apparatus comprises a processor and a non-transitory memory that stores includes instruction causing the processor to perform a method including generating a mask for a current block within a current frame in the video sequence from a residual that is a difference between pixel values of at least two frames other than the current frame, and encoding or decoding the current block by inter-prediction using the mask.

[0005] Another aspect of an apparatus described herein is an apparatus for generating a mask for encoding or decoding a current block of a video signal, the video signal including frames defining a video sequence, the frames having blocks, and the blocks formed of pixels. The apparatus comprises a processor and a non-transitory memory that stores includes instruction causing the processor to perform a method including calculating a residual by

subtracting pixel values within a first frame from pixel values within a second frame, each of the first frame and the second frame located before the current frame within the video sequence, applying a threshold value to pixel values for respective pixel locations within the residual to generate a threshold residual comprising pixels, each pixel within the threshold residual having one of a first value or a second value different from the first value, and expanding at least one of a first area of the threshold residual comprising pixels having the first value or a second area of the threshold residual comprising pixels having the second value to form the mask having a first contiguous portion of pixel locations with the first value and a second contiguous portion of pixel locations with the second value, the first contiguous portion and the second contiguous portion being non-overlapping, and a border between the first contiguous portion and the second contiguous portion extending between two edges of the mask; and wherein the mask has the same pixel dimensions as the current block.

[0006] These and other aspects of this disclosure are described in additional detail in the following detailed description, the appended claims and the accompanying figures.

BRIEF DESCRIPTION OF THE DRAWINGS

[0007] The description herein makes reference to the accompanying drawings described below wherein like reference numerals refer to like parts throughout the several views.

[0008] FIG. 1 is a schematic of a video encoding and decoding system.

[0009] FIG. 2 is a block diagram of an example of a computing device that can implement a transmitting station or a receiving station.

[0010] FIG. 3 is a diagram of a video stream to be encoded and subsequently decoded.

[0011] FIG. 4 is a block diagram of a video compression system according to an aspect of the teachings herein.

[0012] FIG. 5 is a block diagram of a video decompression system according to another aspect of the teachings herein.

[0013] FIG. 6 is a flowchart diagram of a process for encoding or decoding a block by motion vector prediction using a prior frame residual according to one implementation of this disclosure.

[0014] FIG. 7 is a flowchart diagram of a process for generating a mask using a prior frame residual according to one implementation of this disclosure.

[0015] FIGS. 8A-8C are diagrams used to explain the process of FIG. 7.

DETAILED DESCRIPTION

[0016] A video stream may be compressed by a variety of techniques to reduce bandwidth required to transmit or store the video stream. A video stream can be encoded into a bitstream, which can involve compression, and then transmitted to a decoder that can decode or decompress the video stream to prepare it for viewing or further processing. Encoding a video stream can involve parameters that make trade-offs between video quality and bitstream size, where increasing the perceived quality of a decoded video stream can increase the number of bits required to transmit or store the bitstream.

[0017] One technique to achieve superior compression performance exploits spatial and temporal correlation of video signals through spatial and/or motion compensated prediction. Inter-prediction, for example, uses a motion vector to identify a previously-encoded and decoded block that resembles a current block to be encoded. By encoding the motion vector, and the difference between the two blocks, a decoder can re-create the current block. However, objects rarely fall on regular block boundaries within images. This results in coding inefficiencies.

[0018] In contrast, the teachings herein describe the generation and use of an inter-predictor that does not require (e.g., square) blocks so as to better match objects within a frame. This can be implemented by using the residual of a prior frame to create a cliff mask for a block that allows two different motion vectors to be applied to the block. Further details are described after an initial discussion of the environment in which the teachings herein may be used.

[0019] FIG. 1 is a schematic of a video encoding and decoding system 100. A transmitting station 102 can be, for example, a computer having an internal configuration of hardware such as that described in FIG. 2. However, other suitable implementations of the transmitting station 102 are possible. For example, the processing of the transmitting station 102 can be distributed among multiple devices.

[0020] A network 104 can connect the transmitting station 102 and a receiving station 106 for encoding and decoding of the video stream. Specifically, the video stream can be encoded in the transmitting station 102 and the encoded video stream can be decoded in the receiving station 106. The network 104 can be, for example, the Internet. The network 104 can also be a local area network (LAN), wide area network (WAN), virtual private network (VPN), cellular telephone network or any other means of transferring the video stream from the transmitting station 102 to, in this example, the receiving station 106.

[0021] The receiving station 106, in one example, can be a computer having an internal configuration of hardware such as that described in FIG. 2. However, other suitable

implementations of the receiving station 106 are possible. For example, the processing of the receiving station 106 can be distributed among multiple devices.

[0022] Other implementations of the video encoding and decoding system 100 are possible. For example, an implementation can omit the network 104. In another implementation, a video stream can be encoded and then stored for transmission at a later time to the receiving station 106 or any other device having memory. In one implementation, the receiving station 106 receives (e.g., via the network 104, a computer bus, and/or some communication pathway) the encoded video stream and stores the video stream for later decoding. In an example implementation, a real-time transport protocol (RTP) is used for transmission of the encoded video over the network 104. In another implementation, a transport protocol other than RTP may be used, e.g., a Hypertext Transfer Protocol (HTTP)-based video streaming protocol.

[0023] When used in a video conferencing system, for example, the transmitting station 102 and/or the receiving station 106 may include the ability to both encode and decode a video stream as described below. For example, the receiving station 106 could be a video conference participant who receives an encoded video bitstream from a video conference server (e.g., the transmitting station 102) to decode and view and further encodes and transmits its own video bitstream to the video conference server for decoding and viewing by other participants.

[0024] FIG. 2 is a block diagram of an example of a computing device 200 that can implement a transmitting station or a receiving station. For example, the computing device 200 can implement one or both of the transmitting station 102 and the receiving station 106 of FIG. 1. The computing device 200 can be in the form of a computing system including multiple computing devices, or in the form of a single computing device, for example, a mobile phone, a tablet computer, a laptop computer, a notebook computer, a desktop computer, and the like.

[0025] A CPU 202 in the computing device 200 can be a central processing unit. Alternatively, the CPU 202 can be any other type of device, or multiple devices, capable of manipulating or processing information now-existing or hereafter developed. Although the disclosed implementations can be practiced with a single processor as shown, e.g., the CPU 202, advantages in speed and efficiency can be achieved using more than one processor.

[0026] A memory 204 in the computing device 200 can be a read only memory (ROM) device or a random access memory (RAM) device in an implementation. Any other suitable type of storage device can be used as the memory 204. The memory 204 can include code

and data 206 that is accessed by the CPU 202 using a bus 212. The memory 204 can further include an operating system 208 and application programs 210, the application programs 210 including at least one program that permits the CPU 202 to perform the methods described here. For example, the application programs 210 can include applications 1 through N, which further include a video coding application that performs the methods described here. The computing device 200 can also include a secondary storage 214, which can, for example, be a memory card used with a mobile computing device. Because the video communication sessions may contain a significant amount of information, they can be stored in whole or in part in the secondary storage 214 and loaded into the memory 204 as needed for processing.

[0027] The computing device 200 can also include one or more output devices, such as a display 218. The display 218 may be, in one example, a touch sensitive display that combines a display with a touch sensitive element that is operable to sense touch inputs. The display 218 can be coupled to the CPU 202 via the bus 212. Other output devices that permit a user to program or otherwise use the computing device 200 can be provided in addition to or as an alternative to the display 218. When the output device is or includes a display, the display can be implemented in various ways, including by a liquid crystal display (LCD), a cathode-ray tube (CRT) display or light emitting diode (LED) display, such as an organic LED (OLED) display.

[0028] The computing device 200 can also include or be in communication with an image-sensing device 220, for example a camera, or any other image-sensing device 220 now existing or hereafter developed that can sense an image such as the image of a user operating the computing device 200. The image-sensing device 220 can be positioned such that it is directed toward the user operating the computing device 200. In an example, the position and optical axis of the image-sensing device 220 can be configured such that the field of vision includes an area that is directly adjacent to the display 218 and from which the display 218 is visible.

[0029] The computing device 200 can also include or be in communication with a sound-sensing device 222, for example a microphone, or any other sound-sensing device now existing or hereafter developed that can sense sounds near the computing device 200. The sound-sensing device 222 can be positioned such that it is directed toward the user operating the computing device 200 and can be configured to receive sounds, for example, speech or other utterances, made by the user while the user operates the computing device 200.

[0030] Although FIG. 2 depicts the CPU 202 and the memory 204 of the computing device 200 as being integrated into a single unit, other configurations can be utilized. The

operations of the CPU 202 can be distributed across multiple machines (each machine having one or more of processors) that can be coupled directly or across a local area or other network. The memory 204 can be distributed across multiple machines such as a network-based memory or memory in multiple machines performing the operations of the computing device 200. Although depicted here as a single bus, the bus 212 of the computing device 200 can be composed of multiple buses. Further, the secondary storage 214 can be directly coupled to the other components of the computing device 200 or can be accessed via a network and can comprise a single integrated unit such as a memory card or multiple units such as multiple memory cards. The computing device 200 can thus be implemented in a wide variety of configurations.

[0031] FIG. 3 is a diagram of an example of a video stream 300 to be encoded and subsequently decoded. The video stream 300 includes a video sequence 302. At the next level, the video sequence 302 includes a number of adjacent frames 304. While three frames are depicted as the adjacent frames 304, the video sequence 302 can include any number of adjacent frames 304. The adjacent frames 304 can then be further subdivided into individual frames, e.g., a frame 306. At the next level, the frame 306 can be divided into a series of planes or segments 308. The segments 308 can be subsets of frames that permit parallel processing, for example. The segments 308 can also be subsets of frames that can separate the video data into separate colors. For example, the frame 306 of color video data can include a luminance plane and two chrominance planes. The segments 308 may be sampled at different resolutions.

[0032] Whether or not the frame 306 is divided into the segments 308, the frame 306 may be further subdivided into blocks 310, which can contain data corresponding to, for example, 16x16 pixels in frame 306. The blocks 310 can also be arranged to include data from one or more planes of pixel data. The blocks 310 can also be of any other suitable size such as 4x4 pixels, 8x8 pixels, 16x8 pixels, 8x16 pixels, 16x16 pixels or larger. Unless otherwise noted, the terms block and macroblock are used interchangeably herein. The frame 306 may be partitioned according to the teachings herein as discussed in more detail below.

[0033] FIG. 4 is a block diagram of an encoder 400 in accordance with an implementation. The encoder 400 can be implemented, as described above, in the transmitting station 102 such as by providing a computer software program stored in memory, for example, the memory 204. The computer software program can include machine instructions that, when executed by a processor such as the CPU 202, cause the transmitting station 102 to encode video data in the manner described in FIG. 4. The encoder 400 can also

be implemented as specialized hardware included in, for example, the transmitting station 102. The encoder 400 has the following stages to perform the various functions in a forward path (shown by the solid connection lines) to produce an encoded or compressed bitstream 420 using the video stream 300 as input: an intra/inter prediction stage 402, a transform stage 404, a quantization stage 406, and an entropy encoding stage 408. The encoder 400 may also include a reconstruction path (shown by the dotted connection lines) to reconstruct a frame for encoding of future blocks. In FIG. 4, the encoder 400 has the following stages to perform the various functions in the reconstruction path: a dequantization stage 410, an inverse transform stage 412, a reconstruction stage 414, and a loop filtering stage 416. Other structural variations of the encoder 400 can be used to encode the video stream 300.

[0034] When the video stream 300 is presented for encoding, the frame 306 can be processed in units of blocks. At the intra/inter prediction stage 402, a block can be encoded using intra-frame prediction (also called intra prediction) or inter-frame prediction (also called inter prediction or inter-prediction herein). In any case, a prediction block can be formed. In the case of intra-prediction, a prediction block may be formed from samples in the current frame that have been previously encoded and reconstructed. In the case of inter-prediction, a prediction block may be formed from samples in one or more previously constructed reference frames as discussed in more detail below.

[0035] Next, still referring to FIG. 4, the prediction block can be subtracted from the current block at the intra/inter prediction stage 402 to produce a residual block (also called a residual). The transform stage 404 transforms the residual into transform coefficients in, for example, the frequency domain using block-based transforms. Such block-based transforms include, for example, the Discrete Cosine Transform (DCT) and the Asymmetric Discrete Sine Transform (ADST). Other block-based transforms are possible. Further, combinations of different transforms may be applied to a single residual. In one example of application of a transform, the DCT transforms the residual block into the frequency domain where the transform coefficient values are based on spatial frequency. The lowest frequency (DC) coefficient at the top-left of the matrix and the highest frequency coefficient at the bottom-right of the matrix. It is worth noting that the size of a prediction block, and hence the resulting residual block, may be different from the size of the transform block. For example, the prediction block may be split into smaller blocks to which separate transforms are applied.

[0036] The quantization stage 406 converts the transform coefficients into discrete quantum values, which are referred to as quantized transform coefficients, using a quantizer

value or a quantization level. For example, the transform coefficients may be divided by the quantizer value and truncated. The quantized transform coefficients are then entropy encoded by the entropy encoding stage 408. Entropy coding may be performed using any number of techniques, including token and binary trees. The entropy-encoded coefficients, together with other information used to decode the block, which may include for example the type of prediction used, transform type, motion vectors and quantizer value, are then output to the compressed bitstream 420. The compressed bitstream 420 can also be referred to as an encoded video stream or encoded video bitstream, and the terms will be used interchangeably herein.

[0037] The reconstruction path in FIG. 4 (shown by the dotted connection lines) can be used to ensure that both the encoder 400 and a decoder 500 (described below) use the same reference frames to decode the compressed bitstream 420. The reconstruction path performs functions that are similar to functions that take place during the decoding process that are discussed in more detail below, including dequantizing the quantized transform coefficients at the dequantization stage 410 and inverse transforming the dequantized transform coefficients at the inverse transform stage 412 to produce a derivative residual block (also called a derivative residual). At the reconstruction stage 414, the prediction block that was predicted at the intra/inter prediction stage 402 can be added to the derivative residual to create a reconstructed block. The loop filtering stage 416 can be applied to the reconstructed block to reduce distortion such as blocking artifacts.

[0038] Other variations of the encoder 400 can be used to encode the compressed bitstream 420. For example, a non-transform based encoder 400 can quantize the residual signal directly without the transform stage 404 for certain blocks or frames. In another implementation, an encoder 400 can have the quantization stage 406 and the dequantization stage 410 combined into a single stage.

[0039] FIG. 5 is a block diagram of a decoder 500 in accordance with another implementation. The decoder 500 can be implemented in the receiving station 106, for example, by providing a computer software program stored in the memory 204. The computer software program can include machine instructions that, when executed by a processor such as the CPU 202, cause the receiving station 106 to decode video data in the manner described in FIG. 5. The decoder 500 can also be implemented in hardware included in, for example, the transmitting station 102 or the receiving station 106.

[0040] The decoder 500, similar to the reconstruction path of the encoder 400 discussed above, includes in one example the following stages to perform various functions to produce

an output video stream 516 from the compressed bitstream 420: an entropy decoding stage 502, a dequantization stage 504, an inverse transform stage 506, an intra/inter prediction stage 508, a reconstruction stage 510, a loop filtering stage 512 and a deblocking filtering stage 514. Other structural variations of the decoder 500 can be used to decode the compressed bitstream 420.

[0041] When the compressed bitstream 420 is presented for decoding, the data elements within the compressed bitstream 420 can be decoded by the entropy decoding stage 502 to produce a set of quantized transform coefficients. The dequantization stage 504 dequantizes the quantized transform coefficients (e.g., by multiplying the quantized transform coefficients by the quantizer value), and the inverse transform stage 506 inverse transforms the dequantized transform coefficients using the selected transform type to produce a derivative residual that can be identical to that created by the inverse transform stage 412 in the encoder 400. Using header information decoded from the compressed bitstream 420, the decoder 500 can use the intra/inter prediction stage 508 to create the same prediction block as was created in the encoder 400, e.g., at the intra/inter prediction stage 402. At the reconstruction stage 510, the prediction block can be added to the derivative residual to create a reconstructed block. The loop filtering stage 512 can be applied to the reconstructed block to reduce blocking artifacts. Other filtering can be applied to the reconstructed block. In this example, the deblocking filtering stage 514 is applied to the reconstructed block to reduce blocking distortion, and the result is output as an output video stream 516. The output video stream 516 can also be referred to as a decoded video stream, and the terms will be used interchangeably herein.

[0042] Other variations of the decoder 500 can be used to decode the compressed bitstream 420. For example, the decoder 500 can produce the output video stream 516 without the deblocking filtering stage 514.

[0043] As mentioned briefly above, a block may be encoded or decoded by motion vector prediction using a prior frame residual. In general, a mask for the block is generated from a residual calculated between pixels of two frames (e.g., the last two frames before the current frame), and then the block is encoded or decoded by inter-prediction using the mask. In this way, a mask that allows two different motion vectors to be applied to a block can be used to better match objects within an image, improving video compression.

[0044] FIG. 6 is a flowchart diagram of a process 600 for encoding or decoding a block by motion vector prediction using a prior frame residual according to one implementation of this disclosure. The method or process 600 can be implemented in a system such as the

computing device 200 to aid the encoding or decoding of a video stream. The process 600 can be implemented, for example, as a software program that is executed by a computing device such as the transmitting station 102 or the receiving station 106. The software program can include machine-readable instructions that are stored in a memory such as the memory 204 that, when executed by a processor such as the CPU 202, cause the computing device to perform the process 600. The process 600 can also be implemented using hardware in whole or in part. As explained above, some computing devices may have multiple memories and multiple processors, and the steps or operations of the process 600 may in such cases be distributed using different processors and memories. Use of the terms “processor” and “memory” in the singular herein encompasses computing devices that have only one processor or one memory as well as devices having multiple processors or memories that may each be used in the performance of some but not necessarily all recited steps.

[0045] For simplicity of explanation, the process 600 is depicted and described as a series of steps or operations. However, steps and operations in accordance with this disclosure can occur in various orders and/or concurrently. Additionally, steps or operations in accordance with this disclosure may occur with other steps or operations not presented and described herein. Furthermore, not all illustrated steps or operations may be required to implement a method in accordance with the disclosed subject matter. The process 600 may be repeated for each frame of the input signal.

[0046] When the process 600 is an encoding process, the input signal can be, for example, the video stream 300. The input signal can be received by the computing performing the process 600 in any number of ways. For example, the input signal can be captured by the image-sensing device 220 or received from another device through an input connected to the bus 212. The input signal could be retrieved from the secondary storage 214 in another implementation. Other ways of receiving and other sources of the input signal are possible. For example, when the process 600 is a decoding process, the input signal can be an encoded bitstream such as the compressed bitstream 420.

[0047] Using the input signal, a mask is generated for a current block within a current frame in a video sequence from a residual that is a difference between pixel values of two frames other than the current frame at 602. In general, the mask can be generated by calculating a residual between two frames (e.g., two frames before the current frame), thresholding the residual, and cleaning up the thresholded residual. One small tap blur may optionally be added to the mask.

[0048] More specifically, FIG. 7 is a flowchart diagram of a process 700 for generating a mask using a prior frame residual according to one implementation of this disclosure. FIGS. 8A-8C are diagrams used to explain the process 700 of FIG. 7.

[0049] The method or process 700 can be implemented in a system such as the computing device 200 to aid the encoding or decoding of a video stream. The process 700 can be implemented, for example, as a software program that is executed by a computing device such as the transmitting station 102 or the receiving station 106. The software program can include machine-readable instructions that are stored in a memory such as the memory 204 that, when executed by a processor such as the CPU 202, cause the computing device to perform the process 700. The process 700 can also be implemented using hardware in whole or in part. As explained above, some computing devices may have multiple memories and multiple processors, and the steps or operations of the process 700 may in such cases be distributed using different processors and memories.

[0050] For simplicity of explanation, the process 700 is depicted and described as a series of steps or operations. However, steps and operations in accordance with this disclosure can occur in various orders and/or concurrently. Additionally, steps or operations in accordance with this disclosure may occur with other steps or operations not presented and described herein. Furthermore, not all illustrated steps or operations may be required to implement a method in accordance with the disclosed subject matter. The process 700 may be repeated for each block or each frame of the input signal.

[0051] In the process 700, generating the mask includes calculating a residual between two frames at 702. More specifically, the residual may be calculated by subtracting pixel values within a first frame from pixel values within a second frame or vice versa. The first and second frames may be located before the current frame within a video sequence defined by the input signal. The first and second frames may be adjacent frames, but more desirably they are separated by one or more frames within the video sequence and a defined amount of time. The defined amount of time is 200 ms in an example, but other values are possible. The pixel values may represent, for example, the luma components or chroma components of some or all of the pixel locations within the first and second frames.

[0052] In one implementation, the pixel values of pixels within the second frame are subtracted from the pixel values of collocated pixels within the first frame or vice versa. Collocated pixels have the same pixel coordinates within different frames. In one implementation, the pixels within the second frame and the collocated pixels within the first frame are collocated with pixels of the current block. Alternatively, the pixels within the

second frame and the collocated pixels within the first frame are shifted by a motion vector relative to pixels of the current block. In some cases, the pixels in one of the first frame or the second frame may be collocated with the current frame, while the pixels in the other are shifted by a motion vector relative to the current block. Desirably, the pixel values are reconstructed pixel values obtained from the encoding and subsequent decoding process of an encoder, such as that described with respect to FIG. 4.

[0053] When the mask is generated in an encoding process, various techniques may be used to select the two frames. In one particularly simple implementation, the last two adjacent frames before the current frame are used. In another implementation, the last frame before the current frame may be selected, along with the frame most identified as a reference frame for the last frame. In other implementations, other frames may be selected so as to provide a residual for the mask generation process.

[0054] When the mask is generated in a decoding process, the two frames may be discerned from header information within the encoded bitstream as discussed in more detail below.

[0055] The residual can represent the entirety of a frame or only a portion of the frame. That is, the residual can be calculated for the entire dimensions of the frame or for only portions of the frame, such as a block of the frame. An example is shown in FIG. 8A. Therein, all pixel values of a first frame 802 are subtracted from collocated pixel values of a second frame 804. The result is the residual frame (or residual) 806. As can be seen from FIG. 8A, a round object 808, such as a ball, the moon, etc. is moving from a first position in the first frame 802 to a second position in the second frame 804. The residual 806 shows a crescent shape 810 that is the difference between pixel values of the first frame 802 and the second frame 804. In this case, the residual is calculated using the entire area of a frame. However, this calculation or subsequent steps of the process 700 may be performed on a portion of the frames, e.g., a block basis.

[0056] Generating the mask in the process 700 also includes, at 704, applying a threshold to the residual generated at 702. More specifically, the process 700 can include applying a threshold value to pixel values for respective pixel locations within the residual to generate a threshold residual. The threshold residual comprises pixels having the same dimensions as the residual or portion of the residual to which the threshold value is applied. In certain implementations, each pixel within the threshold residual has one of a first value or a second value different from the first value. The threshold value could be a positive value or a negative value, or could define a range of values.

[0057] In one implementation, applying the threshold value includes comparing a pixel value of respective pixel locations within the residual with the threshold value. Then, a first value is assigned to a same pixel location within the threshold residual when the pixel value is less than the threshold value. Alternatively, a second value is assigned to the same pixel location within the threshold residual when the pixel value is greater than the threshold value. For example, the first value may be zero, while the second value may be a maximum pixel value of 256. In one implementation, the second value is a value of 1 when the first value is a value of 0. The threshold value is not required to be any particular value as long as it captures changes in relative movement of an edge of object between the two frames while minimizing the capture of background or minimal changes in the values of the pixels representing the object. One value may be, for example, 75, and other values are possible. Unless a non-island-like residual across two borders of the threshold residual results from the thresholding, the threshold residual is not as useful as a mask. The value may be adjusted, or a decision may be made that the portion of the residual will not generate a mask that is as useful for prediction as another may be.

[0058] Referring to FIG. 8B, for example, a block 812 that is a portion of the residual 806 from FIG. 8A is shown. In the block 812, the edge and hatched areas represent the movement of the round object 808 (e.g., its edge) between the first frame 802 and the second frame 804. Applying the threshold value to the block 812 results in the edge and hatched areas being assigned a value of 1, while other areas are assigned a value of 0. For example, pixel locations within a new block (i.e., the threshold residual) that correspond to pixels within the block 812 having a value within the range of ± 75 are assigned the value of 1, while other pixel locations within the threshold residual that correspond to pixels outside the range are assigned the value of 0. A non-island-like residual is seen across two borders, so the block 812 may generate a useful mask.

[0059] Referring again to FIG. 7, the process 700 for generating a mask may also include modifying the threshold residual. At 706, the threshold residual resulting from applying the threshold to the residual at 704 is modified using, for example, a growth and/or a shrink function on the threshold residual. That is, the threshold residual is cleaned up. In one example, the modification involves recursively applying a grow step only right and down within the threshold residual. In such an implementation, if any neighbor above or to the left is set (i.e., has a value of 1), then the current pixel is set (i.e., is converted to the value of 1). Speed of the recursive grow may be improved by working in larger “chunks” or portions of the threshold residual. For example, if any pixel is set at a bottom of a 4x4 pixel group, then

all pixels within the group are set. Alternatively, a recursive grow function could be applied to grow the area having the other value within the threshold residual. Stated generally, modifying the threshold residual includes applying a growth function to expand an area defined by a minimum number of contiguous pixels having a first value of the two values based on values of pixels adjacent to the area.

[0060] This modification process of 706 described so far results in two generally contiguous areas having separate values. However, there can be areas where the continuity pixels of one value may be broken by individual pixels or groups of pixels of another value. Modifying the threshold residual at 706 may include additional steps to reduce these discontinuities. In one example, modifying the threshold residual includes applying a shrink function to remove an area defined by a maximum number of contiguous pixels having the first value that are surrounded by pixels having the second value of the two values or to remove an area defined by the maximum number of contiguous pixels having the second value that are surrounded by pixels having the first value. By removing the area, it means to change the values so that the first and second values form non-overlapping contiguous regions within a block or frame.

[0061] One example of a mask resulting from modifying a threshold residual is seen in FIG. 8C. The mask 814 is generated by thresholding the block 812 of FIG. 8B and modifying the resulting threshold residual using growth and shrink functions so that from a cliff mask with pixels on one side of a line all have a first value while pixels on the other side of the line all have a second value. While a cliff mask can be used (e.g., just black and white), an optional final step in generating a mask according to the process 700 of FIG. 7 includes applying a blur to a border within the mask. The value of a blur will be discussed in more detail below. At this point, it is noted that the blur results in values about the border that form a smoother transition between the areas. The blur may be one small tap blur formed according to a variety of interpolation techniques. The process 700 ends once the mask is generated.

[0062] Once the mask is generated, it can be used for encoding or decoding a current block. Referring again to FIG. 6, one implementation of encoding or decoding the current block using the mask includes inter-predicting a first prediction block portion at 604, inter-predicting a second prediction block portion at 606, generating a prediction block using the portions at 608, and encoding or decoding the current block using the prediction block at 610.

[0063] When the process 600 is an encoding process, inter-predicting a first prediction block portion at 604 includes performing a first motion search within a reference frame for

pixel values within a first contiguous portion of pixel locations of the current block using the mask. That is, a first motion vector that results in the best match for pixel values within the current block that are collocated with the first contiguous portion of the mask is found. The best match defines the first prediction block portion. Similarly, inter-predicting a second prediction block portion at 606 includes performing a second motion search within a reference frame for pixel values within a second contiguous portion of pixel locations of the current block using the mask. That is, a second motion vector that results in the best match for pixel values within the current block that are collocated with the second contiguous portion of the mask is found. The best match defines the second prediction block portion. Although the use of one reference frame is described, the search may be performed on more than one reference frame, and the best matches need not be in the same reference frame.

[0064] Generating a prediction block using the portions at 608 when the process 600 is an encoding process may include generating the prediction block by combining a result of the first motion search with a result of the second motion search using the mask. This combining may be achieved by combining the pixels values of the best matches into a single prediction block. For example, the prediction block may have pixels at positions within a first portion substantially coincident with first contiguous portion of the mask that have values corresponding to the first prediction block portion and pixels at positions within a second portion substantially coincident with the second continuous portion of the mask that have values corresponding to the second prediction block portion. At the border between the first portion and the second portion of the prediction block, the pixel values are a combination of pixel values in accordance with the blur. For example, the blur could have values between 1 and 0 of 0.75, 0.50 and 0.25 at the border. In this case, pixels at the border region between the two portions include a first set of pixels adjacent to the first portion that have values combining 75% of pixel values of the collocated pixels of the first portion with 25% of pixel values of the collocated pixels of the second portion, a second set of pixels adjacent to the second portion that have values combining 25% of pixel values of the collocated pixels of the first portion with 75% of pixel values of the collocated pixels of the second portion, and a third set of pixels between the first and second border regions that have values combining 50% of pixel values of the collocated pixels of the first portion with 50% of pixel values of the collocated pixels of the second portion. This is a simple example, and other techniques for applying the blurred mask to combine the two sets of pixels to generate a prediction block are within the skill of one in the art based on this disclosure.

[0065] It is worth noting that the mask can be modified for use in the inter-predictions of 604 and 606. That is, for example, the mask can be rotated. This changes the pixels selected for each search from the current block. Performing the motion searches thus comprise performing the first and second motion searches within the reference frame using the mask as rotated—that is, finding the best match for pixels from the current frame that are collocated with each of the separate contiguous portions of the rotated mask. Then, generating the prediction block at 608 similarly uses the mask as rotation to combine the best matches for the portions.

[0066] The mask can also be modified for use in the inter-predictions of 604 and 606 by shifting the mask by a motion vector. Namely, benefits of encoding a portion of the current frame corresponding to the size of the mask (whether all or a portion of the frame) may benefit from adjusting the border between the separate contiguous portions of the mask. The border may be adjusted by, for example, adjusting the pixel values so that the contiguous portion to one side of the mask increases in size and the contiguous portion on the opposite side of the mask decreases in size within the bounds of the mask by one of the motion vectors in a previous (e.g., the last) frame before the current frame. For example, the motion vector used to move the boundary, and hence shift the mask, could be a motion vector of a block of the last frame that is collocated with the current block being predicted.

[0067] Encoding the current block using the prediction block at 610 includes generating a residual for the current block, and encoding the residual into an encoded bitstream with information necessary for decoding the current block. For example, the encoding process could include processing the residual using the transform stage 404, the quantization stage 406, and the entropy encoding stage 408 as described with respect to FIG. 4. The information necessary for decoding the current block may include a mode indicator (sometimes called a flag) that indicates that the current block was encoded using a mask, indicators of which frames were used to generate the mask in the encoder (such as frame IDs), the motion vectors found as a result of the motion searches, the identification of the reference frame, and an indicator of any modification to the mask. For example, if the mask was rotated for use in the encoding process, the bitstream would include such an indication. The information may be included in frame, slice, segment, or block headers, and not all of the information need be transmitted in the same header. Moreover, not all information need be transmitted. For example, if there are no changes to the mask after it is generated (e.g., it is not rotated), there is no need to send an indicator of a modification. Further, if the past two frames are always

used with encoding in this mask mode, there is no need to identify the two frames used within the bitstream. Other modifications are possible.

[0068] As may be understood from this description and the description of the encoding process of FIG. 4, the generation of a single mask may not result in the most efficient coding of a block. Therefore, the processing of FIG. 6 may be incorporated into one or more rate-distortion loops that perform inter-prediction using different masks (or the same masks rotated) so as to find the mask and motion vectors for encoding the current block with the lowest encoding cost (e.g., number of bits to encode).

[0069] When the process 600 is a decoding process, generating a mask from a frame residual at 602 is performed according to FIG. 7. According to one implementation, the frames used to calculate the residual are obtained from the encoded bitstream (e.g., by entropy decoding the header containing the information) when the mask mode is used. Alternatively, the frames may be known by the use of the mask mode. For example, if the prior two adjacent frames to the current frame are always used, there is no need to separately signal the identification of the frames to the decoder.

[0070] A first motion vector for inter-predicting the first prediction block portion at 604 and a second motion vector for inter-predicting the second prediction block portion at 606 may be obtained from a header within the bitstream. Inter-predicting the first prediction block portion at 604 may include generating a first reference block using the first motion vector and applying the mask to the first reference block to generate a first masked reference block (i.e., the first prediction block portion). Similarly, inter-predicting the second prediction block portion at 606 may include generating a second reference block using the second motion vector and applying the mask to the second reference block to generate a second masked reference block (i.e., the second prediction block portion). The prediction block is generated at 608 using the portions in a like manner as in described above with respect to the encoding process.

[0071] Decoding the current block using the prediction block at 610 includes decoding the residual for the current block from the bitstream, and adding the prediction block to the residual. For example, the decoding process could include processing the residual using the entropy decoding stage 502, the dequantization stage 504, and the inverse transform stage 506 as described with respect to FIG. 5. Then, the current block could be reconstructed at reconstructions stage 510 as also described with respect to FIG. 5.

[0072] Pixel prediction is used to reduce the amount of data encoded within a bitstream. One technique is to copy blocks of pixels from prior encoded frames using a motion vector.

In actuality, objects do not often fall on regular block boundaries. According to the teachings herein, a predictor (e.g., a prediction block) that better follows the edge shapes of objects and thus may improve video compression.

[0073] The aspects of encoding and decoding described above illustrate some examples of encoding and decoding techniques. However, it is to be understood that encoding and decoding, as those terms are used in the claims, could mean compression, decompression, transformation, or any other processing or change of data.

[0074] The word “example” is used herein to mean serving as an example, instance, or illustration. Any aspect or design described herein as “example” is not necessarily to be construed as preferred or advantageous over other aspects or designs. Rather, use of the word “example” is intended to present concepts in a concrete fashion. As used in this application, the term “or” is intended to mean an inclusive “or” rather than an exclusive “or”. That is, unless specified otherwise, or clear from context, “X includes A or B” is intended to mean any of the natural inclusive permutations. That is, if X includes A; X includes B; or X includes both A and B, then “X includes A or B” is satisfied under any of the foregoing instances. In addition, the articles “a” and “an” as used in this application and the appended claims should generally be construed to mean “one or more” unless specified otherwise or clear from context to be directed to a singular form. Moreover, use of the term “an implementation” or “one implementation” throughout is not intended to mean the same embodiment or implementation unless described as such.

[0075] Implementations of the transmitting station 102 and/or the receiving station 106 (and the algorithms, methods, instructions, etc., stored thereon and/or executed thereby, including by the encoder 400 and the decoder 500) can be realized in hardware, software, or any combination thereof. The hardware can include, for example, computers, intellectual property (IP) cores, application-specific integrated circuits (ASICs), programmable logic arrays, optical processors, programmable logic controllers, microcode, microcontrollers, servers, microprocessors, digital signal processors or any other suitable circuit. In the claims, the term “processor” should be understood as encompassing any of the foregoing hardware, either singly or in combination. The terms “signal” and “data” are used interchangeably. Further, portions of the transmitting station 102 and the receiving station 106 do not necessarily have to be implemented in the same manner.

[0076] Further, in one aspect, for example, the transmitting station 102 or the receiving station 106 can be implemented using a general purpose computer or general purpose processor with a computer program that, when executed, carries out any of the respective

methods, algorithms and/or instructions described herein. In addition or alternatively, for example, a special purpose computer/processor can be utilized which can contain other hardware for carrying out any of the methods, algorithms, or instructions described herein.

[0077] The transmitting station 102 and the receiving station 106 can, for example, be implemented on computers in a video conferencing system. Alternatively, the transmitting station 102 can be implemented on a server and the receiving station 106 can be implemented on a device separate from the server, such as a hand-held communications device. In this instance, the transmitting station 102 can encode content using an encoder 400 into an encoded video signal and transmit the encoded video signal to the communications device. In turn, the communications device can then decode the encoded video signal using a decoder 500. Alternatively, the communications device can decode content stored locally on the communications device, for example, content that was not transmitted by the transmitting station 102. Other suitable transmitting and receiving implementation schemes are available. For example, the receiving station 106 can be a generally stationary personal computer rather than a portable communications device and/or a device including an encoder 400 may also include a decoder 500.

[0078] Further, all or a portion of implementations of the present invention can take the form of a computer program product accessible from, for example, a tangible computer-usable or computer-readable medium. A computer-usable or computer-readable medium can be any device that can, for example, tangibly contain, store, communicate, or transport the program for use by or in connection with any processor. The medium can be, for example, an electronic, magnetic, optical, electromagnetic, or a semiconductor device. Other suitable mediums are also available.

[0079] The above-described embodiments, implementations and aspects have been described in order to allow easy understanding of the present invention and do not limit the present invention. On the contrary, the invention is intended to cover various modifications and equivalent arrangements included within the scope of the appended claims, which scope is to be accorded the broadest interpretation so as to encompass all such modifications and equivalent structure as is permitted under the law.

What is claimed is:

1. A method for encoding or decoding a video signal using a computing device, the video signal including frames defining a video sequence, the frames having blocks formed of pixels, the method comprising:

generating a mask for a current block within a current frame in the video sequence from a residual that is a difference between pixel values of at least two frames in the video sequence other than the current frame; and

encoding or decoding the current block by inter-prediction using the mask.

2. The method of claim 1, wherein generating the mask comprises:

calculating the residual by subtracting pixel values within a first frame from pixel values within a second frame, the first frame and the second frame being located before the current frame within the video sequence and spaced apart by a plurality of frames and a defined amount of time.

3. The method of either claim 1 or claim 2, wherein the mask has a first contiguous portion of pixel locations and a second contiguous portion of pixel locations, the first contiguous portion and the second contiguous portion being non-overlapping, and a border between the first contiguous portion and the second contiguous portion extending between two edges of the mask, the method further comprising:

performing a first motion search within a reference frame for pixel values within the first contiguous portion of pixel locations of the current block using the mask;

performing a second motion search within the reference frame for pixel values within the second contiguous portion of pixel locations of the current block using the mask; and

generating a prediction block by combining a result of the first motion search with a result of the second motion search using the mask.

4. The method of any one of the preceding claims, wherein generating the mask comprises:

applying a threshold value to pixel values for respective pixel locations within the residual to generate a threshold residual comprising pixels, each pixel within the threshold residual having one of a first value or a second value different from the first value; and

modifying the threshold residual to generate the mask.

5. The method of claim 4, wherein applying the threshold value comprises:
comparing a pixel value of respective pixel locations within the residual with the threshold value;
assigning the first value to a same pixel location within the threshold residual when the pixel value is less than the threshold value; and
assigning the second value to the same pixel location within the threshold residual when the pixel value is greater than the threshold value.
6. The method of either claim 4 or claim 5, wherein modifying the threshold residual comprises:
applying a growth function to expand an area defined by a minimum number of contiguous pixels having the first value based on values of pixels adjacent to the area.
7. The method of claim 6, wherein modifying the threshold residual comprises:
applying a shrink function to remove an area defined by a maximum number of contiguous pixels having the first value that are surrounded by pixels having the second value.
8. The method of any one of claims 4 to 7, wherein the at least two frames include a first frame and a second frame, the method further comprising:
calculating the residual by subtracting the pixel values of pixels within the second frame from the pixel values of collocated pixels within the first frame, and wherein a border of the mask is shifted by a motion vector before encoding or decoding the current block.
9. The method of any one of the preceding claims, further comprising:
rotating the mask;
performing a first motion search within a reference frame using the mask as rotated;
performing a second motion search within the reference frame using the mask as rotated; and combining a block found by the first motion search with a block found by the second motion search using the mask to generate a prediction block for encoding the current block.
10. The method of any one of the preceding claims, further comprising:

receiving an encoded bitstream including a first motion vector, a second motion vector, and an indicator that a mask mode was used to encode the current block; and
after generating the mask:
 applying the mask to a first reference block resulting from the first motion vector to generate a first masked reference block;
 applying the mask to a second reference block resulting from the second motion vector to generate a second masked reference block; and
 combining the first masked reference block and the second masked reference block to generate a prediction block, wherein:
 decoding the current block by inter-prediction using the mask comprises decoding the current block using the prediction block.

11. An apparatus for encoding or decoding a video signal, the video signal including frames defining a video sequence, the frames having blocks, and the blocks formed of pixels, the apparatus arranged to:

 generate a mask for a current block within a current frame in the video sequence from a residual that is a difference between pixel values of at least two frames in the video sequence other than the current frame; and
 encode or decode the current block by inter-prediction using the mask.

12. The apparatus of claim 11, further arranged to generate the mask by:
 calculating the residual as a difference between pixel values within a first frame and pixel values within a second frame in the video sequence;
 comparing a pixel value of respective pixel locations within the residual with a threshold value;
 assigning a first value to a same pixel location within a threshold residual when the pixel value is less than the threshold value;
 assigning a second value to the same pixel location within the threshold residual when the pixel value is greater than the threshold value; and
 expanding at least one of a first area comprising pixels having the first value or a second area comprising pixels having the second value to form the mask having a first contiguous portion of pixel locations with the first value and a second contiguous portion of pixel locations with the second value, the first contiguous portion and the second contiguous portion being non-overlapping, and a border between the first contiguous portion and the

second contiguous portion extending between two edges of the mask.

13. The apparatus of claim 12, further arranged to:

modify the mask by modifying values of pixels about the border between the first contiguous portion and the second contiguous portion before using the mask; and

generate a prediction block including pixels of a first block resulting from a first motion vector associated with the first contiguous portion of pixel locations coincident with the first contiguous portion of pixel locations, pixels of a second block resulting from a second motion vector associated with the second contiguous portion of pixel locations coincident with the second contiguous portion of pixel locations, and combined values of pixels of the first block and the second block coincident with the pixels about the border between the first contiguous portion and the second contiguous portion.

14. The apparatus of any one of claims 11 to 13, further arranged to:

read a first signal within an encoded bitstream associated with the current block of the current frame indicating that the current block was encoded using a mask; wherein generating the mask comprises generating the mask using a residual that is a difference between pixel values of a first frame and of a second frame in the video sequence;

generate a first prediction block within a reference frame using a first motion vector;

generate a second prediction block within a reference frame using a second motion vector;

combine the first prediction block with the second prediction block using the mask to produce a combined prediction block for the current block; and

decode the current block by decoding the current block using the combined prediction block.

15. The apparatus of claim 14, further arranged to:

read at least a second signal within the encoded bitstream associated with the current block indicating the first frame and the second frame;

read a third signal from the encoded bitstream indicating the first motion vector; and

read a fourth signal from the encoded bitstream indicating the second motion vector.

16. The apparatus of either claim 14 or claim 15, wherein the pixel values of the first frame and of the second frame comprise reconstructed pixel values generated by

encoding and decoding the first frame and the second frame.

17. An apparatus for generating a mask for encoding or decoding a current block of a video signal, the video signal including frames defining a video sequence, the frame having blocks formed of pixels, the apparatus being arranged to:

calculate a residual by subtracting pixel values within a first frame from pixel values within a second frame, each of the first frame and the second frame located before a current frame within the video sequence;

apply a threshold value to pixel values for respective pixel locations within the residual to generate a threshold residual comprising pixels, each pixel within the threshold residual having one of a first value or a second value different from the first value; and

expand at least one of a first area of the threshold residual comprising pixels having the first value or a second area of the threshold residual comprising pixels having the second value to form the mask having a first contiguous portion of pixel locations with the first value and a second contiguous portion of pixel locations with the second value, the first contiguous portion and the second contiguous portion being non-overlapping, and a border between the first contiguous portion and the second contiguous portion extending between two edges of the mask; and wherein the mask has the same pixel dimensions as the current block.

18. The apparatus of claim 17, further arranged to:

apply a blur to pixels about the border, the blur defining a region of pixels with values between the first value and the second value.

19. The apparatus of claim 17, further arranged to:

after expanding the at least one of the first area or the second area, apply a shrink function to remove an area defined by a maximum number of contiguous pixels having the first value that are surrounded by pixels having the second value to form the mask.

20. The apparatus of claim 17, further arranged to apply the threshold value by:

comparing a pixel value of respective pixel locations within the residual with the threshold value;

assigning the first value to a same pixel location within the threshold residual when the pixel value is less than the threshold value; and

assigning the second value to the same pixel location within the threshold residual

when the pixel value is greater than the threshold value.

21. Apparatus arranged to carry out the steps of any one of claims 1 to 10.

22. The apparatus of claim 21 comprising a processor and a non-transitory memory storing instructions causing the processor to perform the steps of any one of claims 1 to 10 or to put into effect the apparatus of any one of claims 11 to 20.

1/7




FIG. 1




FIG. 2




FIG. 3




FIG. 4

5/7




FIG. 5

6/7




FIG. 6




FIG. 7




FIG. 8A




FIG. 8B




FIG. 8C

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2016/067792

A. CLASSIFICATION OF SUBJECT MATTER
INV. H04N19/105 H04N19/119 H04N19/137 H04N19/176 H04N19/543
ADD.

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
H04N

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

EPO-Internal, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X A	US 5 177 608 A (OHKI JUNICHI [JP] ET AL) 5 January 1993 (1993-01-05) column 1 - column 18 figures 1-40 -----	1-5,11, 12,17 6-10, 18-22
X A	US 5 969 772 A (SAEKI TAKAYUKI [JP]) 19 October 1999 (1999-10-19) column 1 - column 15 figures 1-23 -----	1-5,11, 12 6-10, 13-22
A	US 5 103 488 A (GEMELLO ROBERTO [IT] ET AL) 7 April 1992 (1992-04-07) column 1 - column 9 figures 1-10 -----	1-22



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

17 February 2017

Date of mailing of the international search report

24/02/2017

Name and mailing address of the ISA/

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040,
Fax: (+31-70) 340-3016

Authorized officer

Ernst, Jens

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2016/067792

Patent document cited in search report		Publication date	Patent family member(s)	Publication date
US 5177608	A	05-01-1993	CA 2051891 A1	21-03-1992
			DE 69126585 D1	24-07-1997
			DE 69126585 T2	02-01-1998
			EP 0476603 A2	25-03-1992
			US 5177608 A	05-01-1993

US 5969772	A	19-10-1999	NONE	

US 5103488	A	07-04-1992	IT 1232109 B	23-01-1992
			US 5103488 A	07-04-1992
