

(51) International Patent Classification:  
*H04N 5/232* (2006.01)(21) International Application Number:  
PCT/US2014/051792(22) International Filing Date:  
20 August 2014 (20.08.2014)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:  
14/038,309 26 September 2013 (26.09.2013) US(71) Applicant: **APPLE INC.** [US/US]; 1 Infinite Loop, Cupertino, California 95014 (US).(72) Inventor: **SILVERSTEIN, D. Amnon**; 1 Infinite Loop, M/S 74-3PA, Cupertino, California 95014 (US).(74) Agent: **MEYERTONS, HOOD, KIVLIN, KOWERT & GOETZEL, P.C.**; Kowert, Robert C., P.O. Box 398, Austin, Texas 78767-0398 (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, 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, 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, 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: IN-STREAM ROLLING SHUTTER COMPENSATION




FIG. 1

(57) Abstract: In-stream rolling shutter compensation may be utilized to modify image data to compensate for detected camera motion. An image processor may perform motion matching on image data received from a camera sensor to determine whether and how the camera is moving. Strips of image data are analyzed to find matching locations between the current image and a previous image by generating graphical profiles for each image strip. The graphical profiles for the current strip are compared to corresponding profiles from the previous image to determine matching locations between the two frames. A motion vector for the strip may be computed based on spatial distances between the match locations of the current image and corresponding match locations of the previous frame. Image data for the current strip may be modified based on the motion vector to compensate for perceived camera motion as it is written out to memory.



---

**Published:**

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

## IN-STREAM ROLLING SHUTTER COMPENSATION

**BACKGROUND**

[0001] Modern digital cameras, such as complementary metal–oxide–semiconductor (CMOS) cameras, frequently employ a method of image acquisition in which each frame is recorded by scanning across the frame row by row from top to bottom (or column by column across the frame) rather than by taking a single snapshot of the entire scene at a single point in time. Thus, not all parts of the image may be recorded at the same time. This method of image acquisition is sometimes called rolling shutter, because the shutter is moving (e.g., rolling) across the image area of the scene.

[0002] Due to the rolling shutter, visual distortions may be introduced and may be visually apparent when the entire image is displayed. In general, any movement of the camera as the image is captured may result in a warping or distortion of the image. Images captured using rolling shutter acquisition may exhibit different types of visual effects or distortions. One type of distortion may be caused by the camera not being held completely still throughout the entire process of scanning a frame. If the camera is moving horizontally, the image gets sheared. If the camera moves vertically, image is stretched or compressed.

[0003] For example, in images captured while a camera is vibrating, such when used from a moving vehicle, a wobble distortion may be created. Video exhibiting this distortion may appear to wobble unnaturally in a jelly-like fashion. Similarly, if a camera moves from side to side, a skew distortion may result in which the image may bend diagonally in one direction or another. If a CMOS camera sensor is used to take an image of a moving object, different areas of the object may appear smeared or to change proportion as compared to other areas of the object.

**SUMMARY**

[0004] In-stream rolling shutter compensation may be utilized to modify image data as it is received from a camera sensor to compensate for detected camera motion. When performing in-stream rolling shutter compensation, an image processor may perform motion matching on image data coming through a camera sensor to determine whether the camera is moving and how the camera is moving (e.g., the direction of motion). As the image data is received, the current image may be analyzed in strips to find matching locations between the current image (e.g., frame) and a previous image. Matches between the two frames may be found by comparing graphical profiles containing summed row and column pixel data for each image strip. For example, luminance data for each pixel in a column (of the current strip) may be combined

together (e.g., summed, averaged, etc.) to compute a column graphical profile while luminance data for pixels in a row may be combined to compute a row graphical profile.

[0005] The horizontal and vertical graphical profiles may be compared to corresponding data from the previous image to determine matching locations between the two frames. A motion vector for the strip may be determined based on the spatial distance between pairs of matching locations (e.g., between the current and previous frames) and may represent camera motion between the two frames. As the image data is written out (e.g., to memory) the current image data may be adjusted (e.g., warped) to compensate for the perceived camera motion based on the computed motion vector.

### **BRIEF DESCRIPTION OF THE DRAWINGS**

[0006] FIG. 1 is a logical block diagram illustrating in-stream rolling shutter compensation, according to one embodiment.

[0007] FIG. 2 is a flowchart illustrating one embodiment of a method for in-stream rolling shutter compensation, as described herein.

[0008] FIG. 3 is a diagram illustrating an image undergoing in-stream rolling shutter compensation, according to one embodiment.

[0009] FIG. 4 is a diagram illustrating image strips and blocks according to one embodiment of in-stream rolling shutter compensation.

[0010] FIG. 5 is a flowchart illustrating one embodiment of a method for in-stream rolling shutter compensation, as described herein.

[0011] FIG. 6 is a diagram illustrating one embodiment of comparing graphical profiles according to one embodiment of in-stream rolling shutter compensation.

[0012] FIG. 7 is a diagram illustrating strip offset vectors according to one embodiment of in-stream rolling shutter compensation.

[0013] FIG. 8 is a flowchart illustrating one embodiment of a method for warping an image strip according to in-stream rolling shutter compensation.

[0014] FIG. 9 is a logical block diagram illustrating a computing device capable of implementing in-stream rolling shutter compensation, as described herein.

[0015] While various embodiments are described herein by way of example for several embodiments and illustrative drawings, those skilled in the art will recognize that embodiments are not limited to the embodiments or drawings described. It should be understood that the drawings and detailed description thereto are not intended to limit the embodiments to the particular form disclosed, but on the contrary, the intention is to cover all modifications,

equivalents and alternatives falling within the spirit and scope of the disclosure. Any headings used herein are for organizational purposes only and are not meant to be used to limit the scope of the description. As used throughout this application, the word "may" is used in a permissive sense (i.e., meaning having the potential to), rather than the mandatory sense (i.e., meaning must). Similarly, the words "include", "including", and "includes" mean including, but not limited to.

### **DETAILED DESCRIPTION OF EMBODIMENTS**

[0016] In-stream rolling shutter compensation may be applied to image data to compensate for camera motion by warping the image data while the image data is received from the image sensor, according to some embodiments. Turning now to Fig. 1, which is a logical block diagram illustrating one embodiment of in-stream rolling shutter compensation. When performing in-stream rolling shutter compensation, an image processor may perform motion matching between current image data and previous image data to determine whether the camera is moving and how the camera is moving (e.g., the direction and extent of motion). For example, image data 110 of FIG. 1 includes arrows representing the direction of example camera motion. Please note that the arrows within image data 110 are merely exemplary and do not represent any particular or actual camera motion.

[0017] In one embodiment, image processor 120 may perform motion matching between image data 110 received from camera image sensor 100 and previous image data. Camera image sensor 100 illustrated in FIG. 1 may represent any of various sorts of image sensors, according to various embodiments. For example, in one embodiment, camera image sensor 120 may represent a CMOS based image sensor. In other embodiments, in-stream rolling shutter compensation may be utilized with image data from sources other than an image sensor. For example, image processor 120 may perform in-stream rolling shutter compensation on image data that was previously captured and stored to memory (or other buffer), storage (e.g., disk drive, CD ROM, etc.), or to another storage mechanism, according to different embodiments. In some embodiments, image processor 120 may receive image data from a camera image sensor, but the image data may also be transferred through one or more other devices/components between the image sensor and the image processor. For example, image processor 120 may receive the image data from a buffer or system memory into which the camera sensor stores the image data as it is captured.

[0018] Additionally, image processor 120 may perform in-stream rolling shutter compensation on image data loaded from a remote source, such as over a local or wide area

network, or the internet. Thus, while described herein in terms of image data received from a camera image sensor, in-stream rolling shutter compensation may be utilized with image data regardless of how that image data is received, stored and/or loaded, according to various embodiments.

5 [0019] Image processor 120 may generate graphical profiles for each image (e.g., current and previous) and may use those graphical profiles to compare the two image to determine one or more match locations from which a motion vector may be computed, as will be described in more detail below. Image processor 120 may then warp the image data 110 based on the  
10 compensated image data 130 may then be written out to memory 140, according to some embodiments. In some embodiments, the image data may be warped as it is written out to memory 140.

[0020] In the following detailed description, numerous specific details are set forth to provide a thorough understanding of claimed subject matter. However, it will be understood by those  
15 skilled in the art that claimed subject matter may be practiced without these specific details. In other instances, methods, apparatuses or systems are not described in detail below because they are known by one of ordinary skill in the art in order not to obscure claimed subject matter.

[0021] Some portions of the detailed description which follow are presented in terms of algorithms or symbolic representations of operations on binary digital signals stored within a  
20 memory of a specific apparatus or special purpose computing device or platform. In the context of this particular specification, the term specific apparatus or the like includes a general purpose computer once it is programmed to perform particular functions pursuant to instructions from program software. Algorithmic descriptions or symbolic representations are examples of techniques used by those of ordinary skill in the signal processing or related arts to convey the  
25 substance of their work to others skilled in the art. An algorithm is here, and is generally, considered to be a self-consistent sequence of operations or similar signal processing leading to a desired result. In this context, operations or processing involve physical manipulation of physical quantities. Typically, although not necessarily, such quantities may take the form of electrical or magnetic signals capable of being stored, transferred, combined, compared or otherwise  
30 manipulated. It has proven convenient at times, principally for reasons of common usage, to refer to such signals as bits, data, values, elements, symbols, characters, terms, numbers, numerals or the like. It should be understood, however, that all of these or similar terms are to be associated with appropriate physical quantities and are merely convenient labels.

[0022] Unless specifically stated otherwise, as apparent from the following discussion, it is appreciated that throughout this specification discussions utilizing terms such as “processing,” “computing,” “calculating,” “determining” or the like refer to actions or processes of a specific apparatus, such as a special purpose computer or a similar special purpose electronic computing device. In the context of this specification, therefore, a special purpose computer or a similar special purpose electronic computing device is capable of manipulating or transforming signals, typically represented as physical electronic or magnetic quantities within memories, registers, or other information storage devices, transmission devices, or display devices of the special purpose computer or similar special purpose electronic computing device.

[0023] An image processor utilizing in-stream rolling shutter compensation may not compare two full image frames (e.g., a complete current frame and a complete previous frame) in order to determine camera motion. For instance, many digital (or CMOS) cameras may not maintain full frame buffers and memory buffer constraints (e.g., timing/performance issues) may prevent the comparison of two full frames using frame buffers. Instead, image processor 120 may analyze an image as multiple image portions, as will be described in more detail below regarding FIG. 2.

[0024] Image processor 120, illustrated in FIG. 1, may represent any of various sorts of image processors or image processing units, which may be implemented in various manners. For instance, in one embodiment image processor 120 may represent an image processing unit implemented in hardware, while in another embodiment image processor 120 may represent an image processing unit implemented in software.

[0025] FIG. 2 is a flowchart illustrating one embodiment of a method for in-stream rolling shutter compensation, as described herein. Image processor 120 may perform in-stream rolling shutter compensation on an image by splitting the image into multiple image strips and by splitting the strip into multiple blocks, according to some embodiments. Thus, as illustrated in block 200 image processor 120 may receive image data (e.g., from image sensor 100) for the next strip of an image and may determine graphical profiles for each block of the image strip, as shown in block 210.

[0026] Image processor 120 may generate graphical profiles for each block based on any of various image characteristic or pixel values. For instance, in one embodiment, luminance values may be combined (e.g., summed or averaged) across a block in columns (of pixels) to generate a vertical graphical profile for the block. For example, image processor 120 may be configured to convert red/blue/green data for pixels into corresponding luminance values and the luminance values for all pixels in a column may be combined together to generate a vertical luminance

profile and multiple column profiles may be combined to generate aggregate profiles, such as a vertical profile for the block.

[0027] Similarly, horizontal graphical profiles may be determined for an image block by combining image characteristics for each row in the strip. Thus, horizontal and vertical graphical profiles may be determined for each block of an image strip. In some embodiments, a single graphical profile may include both a horizontal component and a vertical component, but in other embodiments, horizontal and vertical graphical profiles (e.g., row and column profiles) may be determined separately.

[0028] While in some embodiments, image processor 120 may generate luminance based graphical profiles for each block, in other embodiments graphical profiles may be based on different image characteristics (e.g., different color spaces).

[0029] The exact number and size of blocks used by image processor 120 may vary from embodiment to embodiments. Additionally, in one embodiment, image processor 120 may not use individual blocks at all, but instead may determine a graphical profiles for the entire strip. In some embodiments, the number of blocks may be predetermined, while in other embodiments, the size of each block may be predetermined. In yet other embodiments, the number and/or size of blocks may be determined by image processor 120 one an image by image basis.

[0030] As shown in block 220, image processor 120 may compare the graphical profiles to graphical profiles for a previous image to compute an offset for each block based on the block offsets, as illustrated in block 230. Graphical profiles for the current image strip may be compared to graphical profiles for corresponding strips of a previous image (such as a previous video frame) in order to determine matches points between the two images from which camera motion may be computed, according to one embodiment. For example, luminance profiles for each block of a current strip may be compared to luminance profiles for similarly located blocks in a previous image to find matching locations between the two images, as will be described in more detail regarding FIG. 4.

[0031] In some embodiments, image processor 120 may be configured to determine a horizontal offset for a given block based on comparing a horizontal component of the block's graphical profile to graphical profiles for one or more blocks of the previous image. Similarly, image processor 120 may be configured to determine a vertical offset for a given block based on comparing a vertical component of the graphical profile to graphical profiles for one or more blocks of the previous image.

[0032] In some embodiments, the current and previous luminance profiles may be cross-correlated to locate luminance profile matches. Furthermore, in some embodiments, the spatial

difference between the location of a match point in the current image and a corresponding match point in the previous image may be considered the offset for the block. The offsets of each match point for a block may be used to determine an offset vector for the block.

5 [0033] Additionally, image processor 120 may be configured to store, or otherwise retain, the graphical profiles for the current image for use when processing the next image. In general any suitable manner for storing and/or otherwise retaining the graphical profile information for future use may be used in conjunction with in-stream rolling shutter compensation.

10 [0034] As shown in block 240, image processor 120 may warp the image strip according to the computed offset vector, in one embodiment. For example, if the computed offset vector for a strip indicates motion to the right, the image data may be warped to the left to compensate for the perceived motion, according to one embodiment. The discussion of FIG. 8 below discusses the warping of image data to compensate for camera motion in more detail.

15 [0035] If image processor 120 has completed working on the final strip of the image, as shown by the positive output from decision block 250, image processor 120 may then proceed to begin in-stream rolling shutter compensation on another image as illustrated by block 260, such as to process the next image in a video stream. In some embodiments, image processor 120 may save or store the graphical profiles, block offsets and/or offset vectors generated while processing one image for use when processing the next image.

20 [0036] Instead, if image processor 120 has not finished processing the current image, as indicated by the negative output of decision block 250, image processor 120 may begin processing the next strip of image data for the image, such as by receiving image data for the next strip of an image, as shown by the path from the negative output of block 250 to block 200.

25 [0037] As noted above, in-stream rolling shutter compensation may be applied to image data to compensate for camera motion and image processor 120 may analyze image data for an image as a number of strips. FIG. 3 is a diagram illustrating image data analyzed as strips and blocks for in-stream rolling shutter compensation, according to one embodiment. Image data 300 illustrates an image divided into eight horizontal strips, according to one embodiment. Image data 310 illustrates the same image data divided into multiple vertical strips, according to one embodiment. For instance, while some image sensors may scan a scene using a series of horizontal scans, other camera sensors may scan a scene using a series of vertical scans. Thus, in  
30 some embodiments, image processor 120 may use horizon strips while in other embodiments, image processor may use vertical strips when performing in-stream rolling shutter compensation, as described herein.

[0038] Additionally, as noted above, image processor 120 may divide each strip into multiple blocks and image data 320 illustrates one block for one strip of an image, according to one embodiment. The exact size and nature of the horizontal and vertical divisions use in in-stream rolling shutter compensation may vary from embodiment to embodiment.

5 [0039] FIG. 4 is a diagram illustrating image strips and blocks according to one embodiment of in-stream rolling shutter compensation. As noted above regarding FIG. 3, an image may be divided into image strips and image blocks for analysis during in-stream rolling shutter compensation. Image data 400 illustrates an image divided into multiple horizontal strips and the top-most strip 410 is further illustrated divided into multiple blocks, such as block 420. When  
10 performing in-stream rolling shutter compensation, image processor 120 may analyze the image data for an image received from camera image sensor 100 image block by image block and image strip by image strip.

[0040] As noted above, image processor 120 may analyze the first image strip, such as strip 410, received from camera image sensor 100 prior to receiving the remaining strips. Thus, in-  
15 stream rolling shutter compensation may be performed as the image data is being captured by camera image sensor 100 (e.g., in real-time or near real-time). FIG. 4 also illustrates an example graphical profile 430 for image block 420. For instance, graphical profile 430 may represent a luminance column profile for image block 420, according to one embodiment. Thus, in the example illustrated in FIG. 4, graphical profile 430 may be illustrated as a data graph  
20 representing the combined luminance values across the columns of image block 420.

[0041] While illustrated in FIG. 4 as a data graph representing a luminance based profile as one example, graphical profiles may be computed, generated and/or stored in any of numerous manners, according to different embodiments.

[0042] FIG. 5 is a flowchart illustrating one embodiment of a method for in-stream rolling  
25 shutter compensation, as described herein. As shown in block 500, image processor 120 may generate graphical profiles for the next image block, such as image block 420, in a strip 410, according to one embodiment. As noted above, image processor 120 may generate graphical profiles for each image block of an image strip based on luminance data for the pixels in the image block. Thus, as shown in block 510, image processor 120 may combine luminance data  
30 for pixels in the image block by rows into a horizontal luminance profile for the image block and may also combine luminance data for pixels in the image block by columns into a vertical luminance profile for the image block, as illustrated by block 520.

[0043] Luminance data (or other graphical information) for pixels in the image block may be combined in any of various manners in order to generate graphical profiles, according to different

embodiments. For example, in one embodiment, luminance values for the pixels in a row (or column) may be averaged together to generate a graphical profile, while in other embodiments, the luminance values may be added together to generate a graphical profile. Additionally, in some embodiments, one or more filters may be applied to the data when generating graphical profiles. For example, when generating a luminance based graphical profile, a highpass filter may be used. The use of such filters may aid the finding of profile matches, as discussed below.

[0044] As noted previously, the spatial offset between two matching luminance profile locations (one in the current image and one in the previous image) may represent camera motion between the two images. The luminance profiles may be compared to identify similar locations within each image (e.g., the current and previous) in order to determine the camera motion between the two images.

[0045] Thus, as shown in block 530, image processor 120 may compute the sum of differences for possible matches between the current image block and corresponding image blocks of a previous image using the graphical profiles and may determine a best fit match and use the best fit match as a block offset, as shown in block 540. For example, by comparing the sum of differences for multiple (or all) possible matches between the current block and corresponding blocks of the previous image, potential matches locations may be found between the two images. Locations in each image that have similar graphical profiles may be considered matches or potential matches for the two images.

[0046] Luminance profile matches (or potential matches) may be determined using 2 one-dimensional matches, rather than using 1 two-dimensional match, in some embodiments. For example, using 2 one-dimensional matches may result in lower computational requirements than a single two-dimensional match. In other words, for each block, a horizontal (e.g., left/right) match and a vertical (e.g., up/down) match may be found separately (e.g., 2 one-dimensional matches).

[0047] In order to find matches, the luminance profile for a current block may be compared to luminance profiles for several blocks for corresponding locations (and nearby locations) from the previous image or frame. Any potential matches may be categorized and compared to other potential matches in order to determine a best fit, is possible. For instance, the sum of absolute differences for a potential match may be used to categorize the potential match. If the sum of absolute differences for a match is small, the match may be considered a good fit. If the sum of absolute differences for a potential match is small, but the sum of absolute differences for other areas around the potential match (e.g., for other potential matches) is large, the fit for the potential match may be considered unique.

[0048] Thus, in some embodiments, image processor 120 may be configured to determine that a possible (or potential) match is a good fit if the sum of absolute differences between the pixel location of the match within the current block and the corresponding pixel location of the match within a block of the previous image is smaller than a predetermined threshold.

5 Similarly, image processor may, in some embodiments, be configured to determine that a possible match is a unique fit if the sum of absolute differences (e.g., between the pixel location of the match within the current block and the corresponding pixel location of the match within a block of the previous image) is smaller, by a predetermined threshold, than the respective sums of absolute differences for other possible matches. Furthermore, image  
10 processor may be configured to determine the block offset based on the offsets for one or more matches that were each determined to be either a good fit or a unique fits, according to some embodiments.

[0049] In other words, if there is a position within a block for which a sharp minimum of the sum of absolute differences occurs, that position may represent a unique fit between the two  
15 images. If match may be considered unique, that match location may represent a logical reference point between the two images (e.g., may represent the same object or location in the scene being captured).

[0050] For example, FIG. 6 is a diagram illustrating one embodiment of comparing graphical profiles according to one embodiment of in-stream rolling shutter compensation. Graphical  
20 profiles 610 and 630 may represent luminance based graphical profiles to which highpass filters have been applied. According to the example embodiment shown in FIG. 6, graphical profile 610 may correspond to image block 600a while graphical profile 630 may correspond to image block 600b. In one embodiment, the sum of differences may be represented by the line 620 in FIG. 6. As illustrated by the highest point of line 620, a potential match may exist at that  
25 location within the relevant image blocks.

[0051] When searching for potential matches using graphical profiles as part of in-stream rolling shutter compensation, image processor 120 may not search over a large (or relatively large) area of the image because it may be assumed that the camera is not moving very quickly between image frames. For instance, if it is assumed that the camera is not being moved to  
30 intentionally create a pan (e.g., if the camera is being held still), then it may be assumed that the camera will not move much between image frames.

[0052] Additionally, in some embodiments, image processor 120 may be configured to use an estimated or predicted location in which to search for potential matches using graphical profiles. In one embodiment, image processor 120 may be configured to use an offset vector recently

determined (e.g., for a recently analyzed block and/or strip) to estimate the location of potential matches for the current block/strip. For example, if image processor 120 is currently processing a strip near row 1000 in the overall image and a recently determined offset vector indicates a distance (in terms of dx, dy pixel) of -123, 456 pixels, image processor 120 may begin searching for a potential match for the current block/strip at a location (in the previous image) that is -123, 456 pixels from the strip's location in the current image (e.g., row 1456). Thus, rather than having to search over a large area (or even the entire previous image) for potential matches, image processor 120 may be able to search over a smaller area and therefore load a smaller number of graphical profiles for the previous image, according to some embodiments.

**[0053]** An image processor performing in-stream rolling shutter compensation may then use the luminance profile matches to determine the sub-block motion match in terms of x/y image coordinates. Additionally, multiple matches (e.g., fits) occurring in multiple blocks may be used to determine an aggregate (or average) motion vector for the entire image strip. In some embodiments, it may be assumed that a single strip should have a mostly uniform motion across the entire strip if the perceived motion is caused by camera motion (as opposed to scene motion).

**[0054]** Returning now to FIG. 5, if image processor 120 has finished processing all the block for the current strip, as indicated by the positive output from decision block 550, image processor 120 may then combine the individual block offsets into a strip offset vector, as illustrated by block 560. Thus, a motion vector may be derived based on the offsets from one or more match points (in terms of pixel locations) between the two images may be used to warp the image data to compensate for the determined camera motion. The determined motion vector may include only horizontal and vertical components, without any rotation component, according to some embodiments.

**[0055]** For example, FIG. 7 is a diagram illustrating strip offset vectors according to one embodiment of in-stream rolling shutter compensation. Diagram 700 illustrates an example set of block offsets (or block vectors) for an example image divided into four strips, according to one embodiment. The example offsets illustrated in diagram 700 are represented as arrows, the direction of which may represent the direction of camera motion and the length of which may represent the magnitude of that motion, according to the example embodiment shown in FIG. 7. Thus, the example camera motion illustrated by diagram 700 may represent a motion that begins going horizontally to the left and gradually (over the course of the frame) become a slower movement diagonally down to the left.

**[0056]** As noted above, image processor 120 may compute a single strip vector based on multiple individual block offsets. In one embodiment, image processor 120 may utilize vector

arithmetic to combine the individual block offsets (e.g., as vectors) of a strip to determine a strip motion vector. In some embodiments, the individual block offsets may be averaged together, while in other embodiments, the strip motion vector may be mean of the individual block offsets. For instance, arrows 710, 720, 730 and 740 may represented the computed strip motion vectors determined for each respective strip and based on the corresponding block offsets illustrated in diagram 700.

[0057] Additionally, one or more individual block offsets may be ignored, such as if they represent motion outside of some tolerance of the other block offsets for the strip. For example, in diagram 700 individual block offset 750 may be the result of actual motion in the scene being captured rather than the camera motion causing the other block offsets and therefore block offset 750 may not be considered when computing the strip motion vector for that strip, such as strip vector 730.

[0058] If no motion is determined for a strip (e.g., no match, no match considered good enough, or no motion detected for a match), the motion vector for the previous strip may be used to determine a corrective warp applied to the current strip, according to some embodiments. Thus, even if no motion may be detected for the current strip, such as because the relevant portion of the actual scene does not include a “matchable” point, the previous motion vector is continued (while possibly being modified). For instance, if no motion vector is determined for the top-most strip of an image frame, the motion used for the bottom-most strip of the previous frame may be used, in some embodiments.

[0059] Returning again to FIG. 5, if, however, there are more blocks to process for the current strip, as indicated by the negative output from decision block 550, image processor 120 may proceed to begin processing the next strip as indicated by the path from the negative output from decision block 550 to block 500.

[0060] The following set of example pseudo code illustrates one embodiment of a function implementing in-stream rolling shutter compensation including the computing of a motion vector for a strip.

```
function [warp_vectors, global_vector, block_vectors, old_vector] =  
    FindMotionBlocks(Y1,Y2,chunks,strips,max_move,old_vector, vblank)  
  
    % Takes two images (Y1 and Y2) with rolling shutter and finds a strip-wise  
    % estimate of the motion  
  
    % It starts with cutting each strip into chunks and it boils the chunks  
    % down to a global estimate for the motion of the entire strip  
  
    % Max move is the size of the search area in pixels (up, down, left and  
    % right)  
  
    % The old vector [x, y, x quality, y quality] is from the bottom strip in the
```

```

% previous image. The motion at the end of the last frame is used to start
% the estimate for the top of the next frame

5   if exist('chunks','var')~=1 chunks=5; end
   if exist('strips','var')~=1 strips=8; end
   if exist('max_move','var')~=1 max_move=8; end
   if exist('old_vector','var')~=1 old_vector=[0 0 0 0]; end

10  [h,w]=size(Y1);
   x_size=(w-(2*max_move))/chunks;
   y_size=(h-(2*max_move))/chunks;

   vectors=zeros(chunks,strips,2); % Best motion match, horizontal and vertical
   fits=zeros(chunks,strips,2);    % Quality of match in horizontal and vertical
15  centers=zeros(chunks,strips,2); % Center of target used for the match

   for x=0:(chunks-1)
       for y=0:(strips-1)

20         % Center of the blocks. Blocks have an even grid of centers across
           % the image with a border of max_move pixels, to provide support
           % for the motion match.

           xc= ((x+.5)/chunks) * (w-(2*max_move)-1) + max_move+1;
25         yc= ((y+.5)/strips) * (h-(2*max_move)-1) + max_move+1;

           block1=Y1( round((-y_size/2):(y_size/2))+yc,round((-x_size/2):(x_size/2))+xc);
           block2=Y2( round((-y_size/2-max_move):(y_size/2+max_move))+yc,round((-x_size/2-max_move):(x_size/2+max_move))+xc);
30         % BlockMatch takes two luminance image blocks
           % and it finds the motion vectors
           [vector,fit]=BlockMatch(block1,block2,max_move);
35         vectors(y+1,x+1,:)=vector;
           fits(y+1,x+1,:)=fit;
           centers(y+1,x+1,:)= [xc yc];
       end
   end
40   % Sort out the motion vectors by fit quality

   % A good fit is better than half median

45   h_fits=fits(:, :, 1);
   good_h=median(h_fits(:))/2;
   v_fits=fits(:, :, 2);
   good_v=median(v_fits(:))/2;
   good_fits(:, :, 1)=h_fits>good_h;
50   good_fits(:, :, 2)=v_fits>good_v;

   % Strip vectors have six elements
   % x,y,dx,dy,qx,qy
   % There is also a strip for the previous frame and for the next frame.
55   strip_vectors=zeros(size(vectors,1)+2,6);
   strip_vectors(1,3:6)=old_vector;

   for strip=1:size(fits,1)
       for way=1:2
90         good_vectors=vectors(strip, :, way);
           good_vectors=good_vectors(good_fits(strip, :, way));
           if size(good_vectors(:),1)>0
               % If we have at least one good vector, use the median
           65         v=median(good_vectors(:));
               else
                   v=old_vector(way);
               end
           good_count=mean(good_fits(strip, :, way));
           if (good_count<.5 && strip_vectors(strip,way+4)>0)
               % This strip doesn't have enough good motion tracking
               % Just assume constant motion if we have a good previous vector
           70         v=strip_vectors(strip,way+2);
               end
       end
   end

```

```

        strip_vectors(strip+1,way+2)=v; % best guess of motion vector for the strip
    end
    strip_vectors(strip+1,1)=size(Y1,2)/2; % Strip center is just the middle width of
    image
5    strip_vectors(strip+1,2)=centers(strip,1,2); % Center of the strip
    strip_vectors(strip+1,way+4)=good_count;
end

% Fill in the first vector from the old frame.
10    strip_vectors(1,1)=strip_vectors(2,1);
    strip_vectors(1,2)=strip_vectors(end-1,2)-h*(1+vblank); % lines back in time of the
    bottom
    strip_vectors(end,:)=strip_vectors(end-1,:);
15    strip_vectors(end,2)=h+1;

    old_vector=strip_vectors(end-1,3:6);

% Median or mean? Depends on how confident we are
20    global_vector=[mean(strip_vectors(:,3)) mean(strip_vectors(:,4))];

    warp_vectors=strip_vectors;
    % Already filled in. warp_vectors(2:(end-1),5:6)=mean(good_fits,2);
    % Some estimate of how believable the vectors are.
25    block_vectors=strip_vectors;

```

[0061] The following set of example pseudo code illustrates one embodiment of a function implementing in-stream rolling shutter compensation including the computing of a motion vector for a strip. The BlockMatch function illustrated below is called by the previous function

30 illustrated in the example source code above.

```

function [motion,fit_quality]=BlockMatch(target, neighborhood, max_move)
    if exist('max_move','var')~=1 max_move=8; end

35    % Takes two luminance image blocks and it finds the motion vectors for rolling
    % shutter etc.

    % It uses strip-wise row and column sums

40    % Input should be top-to-bottom scan order

    % Returns a motion vector dx,dy for the best guess of motion
    % Returns a fit value that indicates the quality of the motion match

45    % Low-pass filter
    lp_size=8;
    k=normpdf(-lp_size:lp_size,0,lp_size/2);
    k=k./sum(k);

50    % Two blocks, two dimensions, high pass sums. Blocks can be rectangles
    % of different sizes.
    blocks={target,neighborhood};
    blocks_hp={{[],[]},{[],[]}};

55    for b=1:2 % Target and neighborhood
        block=blocks{b};
        for way=1:2 % Horizontal and vertical
            oversize=size(block,way)-size(target,way);

60            if way==2
                block=block';
            end

            % Instead of summing the entire support, sum just the overlap
            % part with wings that go out further
65            if oversize>0
                strip=block((oversize/2):(oversize/2)+size(target,way)),:);
            end
        end
    end

```

```

else
    strip=block;
end
5      block_sum=sum(strip)';

    % Pad support with copies
    for i=2:size(block_sum,1)
        if isnan(block_sum(i))
10         block_sum(i)=block_sum(i-1);
        end
    end

    % Pad support with copies
    for i=(size(block_sum,1)-1):-1:1
        if isnan(block_sum(i))
15         block_sum(i)=block_sum(i+1);
        end
    end

    % Support LP with mirror.
    block_low_pass=conv([flipud(block_sum);block_sum;flipud(block_sum)],k,'same')
    ;
    block_low_pass=block_low_pass((size(block_sum,1)+1):(size(block_sum,1)*2));
25     block_high_pass=block_sum-block_low_pass;
    blocks_hp{b}{way}=block_high_pass;
end
end

30     motion=[0,0];
    fit_quality=[0,0];

    % Horizontal then vertical match
    for way=1:2
35         target_profile=blocks_hp{1}{way};
        neighborhood_profile=blocks_hp{2}{way};

        %fit=conv(neighborhood_profile,flipud(target_profile),'valid');
        fit=SlidingSAD(target_profile,neighborhood_profile);
40         center=round(size(fit,1)/2);
        area_of_interest=fit((center-max_move):(center+max_move),:);
        [peak_height,peak]=min(area_of_interest);
        motion(way) = peak - round(size(area_of_interest,1)/2);

45         % A test of fit uniqueness.
        fit_min=max(peak-round(size(area_of_interest,1)/10),1);
        fit_max=min(peak+round(size(area_of_interest,1)/10),size(area_of_interest,1));
        other_fits=area_of_interest([1:fit_min fit_max:end]);

50         fit_quality(way)=min(other_fits)-peak_height;
    end
end

```

[0062] While the above example code snippets uses particular function, method and variable names, other embodiments may utilize different functions, methods and/or variables and may perform the same, more or fewer functions when implementing in-stream rolling shutter compensation, as described herein.

[0063] FIG. 8 is a flowchart illustrating one embodiment of a method for warping an image strip according to in-stream rolling shutter compensation. As noted previously image processor 120 may perform a warp on the image data for an individual image strip to compensate for the determined camera motion. Thus, image processor 120 may begin to warp the strip according to the determined offset vector according to one embodiment, as shown in block 800.

[0064] The exact nature of the warp may be determined based on the motion vector determined for the strip. Such a warp may be performed when writing the image data out to memory (or another component of the digital camera). Image processor 120 may warp the image data for the current strip as that image data is written out to memory and before beginning the analysis of the next image strip, according to some embodiments. The warp may include stretching or shearing of the image (in both the vertical and horizontal directions), according to some embodiments. The corrective (e.g., compensative) warp may be considered separable, according to some embodiments. That is, each line in the image data, may be shifted to the right or left horizontally and may also be stretched or compressed vertically.

[0065] The corrective (e.g., compensative) warp may include a horizontal shift of the image data within the strip, such as by skipping a number of pixels at one end of each line in the strip while padding the other end of the line (e.g., with zeros). For instance, if there is a rightward horizontal component in the offset vector, as indicated by the positive output from decision block 810, then image processor 120 may omit one or more left-most columns from the strip to compensate for the rightward horizontal vector component, as shown in block 820. Alternatively, if there is a leftward horizontal component in the offset vector, as indicated by the positive output from decision block 830, image processor 120 may omit one or more right-most columns to compensate for the leftward horizontal vector component, as illustrated by block 840.

[0066] The warp may also include a vertical stretch or compression to compensate for perceived vertical motion. In one embodiment, the image may be stretched vertically by writing a line multiple times while lines may be skipped in order to compress the image vertically. For instance, if there is a downward vertical component in the offset vector, as indicated by the positive output from decision block 850, image processor 120 may combine or merge two or more rows in the strip to compensate for the downward vertical vector component, as shown in block 860. Alternatively, if there is an upward vertical component in the offset vector, as indicated by the positive output of decision block 870, image processor 120 may expand, duplicate, or resample one or more rows to compensate for the upward vertical vector component, as shown in block 880.

[0067] Thus, image processor 120 may be configured to output additional or fewer rows of image data that being received for a particular strip to compensate for detected camera motion. Image processor 120 may combine or merge multiple rows in any of various manners, according to different embodiments. For example, in one embodiment, image processor 120 may be configured to average or interpolate the image data of the rows in order to combine or merge two or more rows. In another example, image processor 120 may resample one or more rows of

image data to obtain fewer rows as a way of combining or merging the rows. Image processor 120 may also resample one or more rows to output additional rows. In general any suitable method for combining, merging, expanding duplicating, and/or resampling the image data may be used to warp (or otherwise modify) image data rows to compensate for detected camera motion, according to some embodiments.

[0068] Thus, image processor 120 may perform vertical stretching or compression by generating fewer or additional lines of image data than in the original (e.g., as received from the camera image sensor). Additionally, image processor 120 may be configured to adjust the data and write it out faster than it is received, according to some embodiments. Therefore, in some embodiments, image processor 120 may be configured to perform in-stream rolling shutter compensation quickly enough to be done as the image data is received from the camera image sensor (e.g., in real-time or near real-time).

[0069] In general, whatever motion was determined (and represented by computed motion vectors) may be compensated for as the image data is written out, according to some embodiments. In some embodiments, no rotation compensation may be applied when warping the image data to compensate for detected camera motion. Image processor 120 may then continue processing the image by starting on the next image strip, as illustrated by block 890.

[0070] Please note that while described as separate functions or steps for each of discussion, the warping of a line of image data may be performed in such a way as to incorporate both a vertical and a horizontal compensation. For instance, in one embodiment, if the computed motion vector for a strip includes both a horizontal component and a vertical component, a warp may be applied to compensate for both motion components. Thus, a line of image data may be duplicated, such as to compensate for a vertical vector component, and may also have one or more pixels omitted from one line end, such as to compensate for a horizontal vector component, according to one embodiment.

[0071] When performing in-stream rolling shutter compensation, one or more constraints may be used to limit the amount of motion compensation may vary from one strip to the next. For example, a damping factor may be applied to help ensure that the corrective warp is smooth (e.g., no abrupt camera motion may be apparent).

[0072] In general, only a small number of line buffers may be needed to perform in-stream rolling shutter compensation. In one embodiment, the number of line buffers may be equal to the number of columns in the strip plus one extra (e.g., for combining two lines – one in the strip and one adjacent to the strip). For example, if a strip is 8 lines high (in terms of pixels) 9 line buffers may be used when perform in-stream rolling shutter compensation, as described herein.

[0073] According to some embodiments, in-stream rolling shutter compensation may be performed while the image is being captured (e.g., in real-time or near real-time) rather than as a post-processing step (e.g., after the image data has been saved to system memory). However, the methods and/or techniques for in-stream rolling shutter compensation described herein may apply to, and be used during, post-processing of an image.

[0074] The techniques described herein regarding in-stream rolling shutter compensation do not utilize gyroscopic data from camera hardware that includes a gyroscope. According to some embodiments, just the image data alone is used to determine the camera motion. In general, use of gyroscopic data may be too slow for use when in-stream rolling shutter compensation, but in some embodiments gyroscopic data could be used to verify that the determined motion (e.g., which is being compensated for) is actually camera motion and not scene motion. For example, in one embodiment, image processor 120 may be configured to check gyroscopic data after processing multiple image frames using in-stream rolling shutter compensation. If the motion is determined to not be actual camera motion according to the gyroscopic data, image processor 120 may be configured to stop performing in-stream rolling shutter compensation.

[0075] Additionally, the techniques described herein for in-stream rolling shutter compensation may, in some embodiments, be considered a fall back mode to hardware based Graphical Processing Unit (GPU) motion correction. Alternatively, image processor 120 may be configured to activate a GPU to reprocess frames if checking the gyroscopic data determines that scene motion was incorrectly detected as camera motion.

[0076] The techniques described herein for in-stream rolling shutter compensation may be implemented on any of a wide variety of computing devices. Turning now to Fig. 9, which illustrates one embodiment of system configured to implement in-stream rolling shutter compensation, as described herein. In general, the computer system illustrated in FIG. 9 may be any of various types of computing device, including, but not limited to, a personal computer system, desktop computer, laptop or notebook computer, mainframe computer system, handheld computing device, cellular phone, smartphone, mobile phone, workstation, network computer, a consumer device, application server, storage device, intelligent display, a peripheral device such as a switch, modem, router, etc., or in general any type of computing device. According to one example embodiment, the system illustrated in FIG. 9 may represent a system-on-a-chip (SOC). As implied by the name, the components of the SOC 900 may be integrated onto a single semiconductor substrate as an integrated circuit "chip." In some embodiments, the components may be implemented on two or more discrete chips in a system. The SOC 900 will be used as an example herein.

[0077] In the illustrated embodiment, the components of the SOC 900 include a central processing unit (CPU) complex 920, on-chip peripheral components 940A-940B (more briefly, "peripherals"), a memory controller (MC) 930, a communication fabric 910, and an image processor 120 (which may itself be considered a peripheral component). SOC 900 may also be coupled to additional components, such as to memory 1800 and camera image sensor 100. The components 920, 930, 940A-940B, and 120 may all be coupled to the communication fabric 910. The memory controller 930 may be coupled to the memory 1800 during use, and the peripheral 940B may be coupled to an external interface 1900 during use. Additionally, image processor 120 may be coupled to camera image sensor 100 during use. In the illustrated embodiment, the CPU complex 920 includes one or more processors 924 and a level two (L2) cache 922.

[0078] The peripherals 940A-940B may be any set of additional hardware functionality included in the SOC 900. For example, the peripherals 940A-940B may include display controllers configured to display video data on one or more display devices, graphics processing units (GPUs), video encoder/decoders, scalers, rotators, blenders, etc. Image processor 120 may, in some embodiments, be part of another video peripheral such as an image signal processor configured to process image capture data from camera image sensor 100 (or other image sensor). The peripherals may also include audio peripherals such as microphones, speakers, interfaces to microphones and speakers, audio processors, digital signal processors, mixers, etc. The peripherals may include peripheral interface controllers for various interfaces 1900 external to the SOC 900 (e.g. the peripheral 940B) including interfaces such as Universal Serial Bus (USB), peripheral component interconnect (PCI) including PCI Express (PCIe), serial and parallel ports, etc. The peripherals may further include networking peripherals such as media access controllers (MACs). In general, any set of hardware may be included, according to various embodiments.

[0079] More particularly in Fig. 9, SOC 900 may include at least one instance of an image processor component, for example image processor 120 that may be configured to implement in-stream rolling shutter compensation, as described herein.

[0080] The CPU complex 920 may include one or more CPU processors 924 that serve as the CPU of the SOC 900. The CPU of the system may include the processor(s) that execute the main control software of the system, such as an operating system. Generally, software executed by the CPU during use may control the other components of the system to realize the desired functionality of the system. The processors 924 may also execute other software, such as application programs. The application programs may provide user functionality, and may rely on the operating system for lower level device control. Accordingly, the processors 924 may also be referred to as application processors. The CPU complex 920 may further include other hardware

such as the L2 cache 922 and/or interface to the other components of the system (e.g. an interface to the communication fabric 910).

[0081] Generally, a processor may include any circuitry and/or microcode configured to execute instructions defined in an instruction set architecture implemented by the processor. The instructions and data operated on by the processors in response to executing the instructions may generally be stored in the memory 1800, although certain instructions may be defined for direct processor access to peripherals as well. Processors may encompass processor cores implemented on an integrated circuit with other components as a system on a chip (SOC 900) or other levels of integration. Processors may further encompass discrete microprocessors, processor cores and/or microprocessors integrated into multichip module implementations, processors implemented as multiple integrated circuits, etc.

[0082] The memory controller 930 may generally include the circuitry for receiving memory operations from the other components of the SOC 900 and for accessing the memory 1800 to complete the memory operations. The memory controller 930 may be configured to access any type of memory 1800. For example, the memory 1800 may be static random access memory (SRAM), dynamic RAM (DRAM) such as synchronous DRAM (SDRAM) including double data rate (DDR, DDR2, DDR3, etc.) DRAM. Low power/mobile versions of the DDR DRAM may be supported (e.g. LPDDR, mDDR, etc.). The memory controller 930 may include queues for memory operations, for ordering (and potentially reordering) the operations and presenting the operations to the memory 1800. The memory controller 930 may further include data buffers to store write data awaiting write to memory and read data awaiting return to the source of the memory operation. In some embodiments, the memory controller 930 may include a memory cache to store recently accessed memory data. In SOC implementations, for example, the memory cache may reduce power consumption in the SOC by avoiding reaccess of data from the memory 1800 if it is expected to be accessed again soon. In some cases, the memory cache may also be referred to as a system cache, as opposed to private caches such as the L2 cache 1122 or caches in the processors 924, which serve only certain components. Additionally, in some embodiments, a system cache need not be located within the memory controller 930.

[0083] In an embodiment, the memory 1800 may be packaged with the SOC 900 in a chip-on-chip or package-on-package configuration. A multichip module configuration of the SOC 900 and the memory 1800 may be used as well. Such configurations may be relatively more secure (in terms of data observability) than transmissions to other components in the system (e.g. to the end points 16A-16B). Accordingly, protected data may reside in the memory 1800

unencrypted, whereas the protected data may be encrypted for exchange between the SOC 900 and external endpoints.

[0084] The communication fabric 910 may be any communication interconnect and protocol for communicating among the components of the SOC 900. The communication fabric 910 may be bus-based, including shared bus configurations, cross bar configurations, and hierarchical buses with bridges. The communication fabric 910 may also be packet-based, and may be hierarchical with bridges, cross bar, point-to-point, or other interconnects.

[0085] It is noted that the number of components of the SOC 900 (and the number of subcomponents for those shown in Fig. 9, such as within the CPU complex 920) may vary from embodiment to embodiment. There may be more or fewer of each component/subcomponent than the number shown in Fig. 9.

[0086] In some embodiments, the methods described herein may be implemented by a computer program product, or software. In some embodiments a non-transitory, computer-readable storage medium may have stored thereon instructions which may be used to program a computer system (or other electronic devices) to perform some or all of the techniques described herein. A computer-readable storage medium may include any mechanism for storing information in a form (e.g., software, processing application) readable by a machine (e.g., a computer). The machine-readable medium may include, but is not limited to, magnetic storage medium (e.g., floppy diskette); optical storage medium (e.g., CD-ROM); magneto-optical storage medium; read only memory (ROM); random access memory (RAM); erasable programmable memory (e.g., EPROM and EEPROM); flash memory; electrical, or other types of medium suitable for storing program instructions. In addition, program instructions may be communicated using optical, acoustical or other form of propagated signal (e.g., carrier waves, infrared signals, digital signals, etc.).

[0087] A computer system 900 may include a processor unit 920 (possibly including multiple processors, a single-threaded processor, a multi-threaded processor, a multi-core processor, etc.) which may be configured to execute one or more modules, such as image processor 120, which may be present within program instructions stored in memory 1800 of the same computer system or may be present within program instructions stored within a memory of another computer system similar to or different from computer system 900. In some embodiments, the various components configured to implement in-stream rolling shutter compensation may be part of different computer systems working together.

[0088] System memory 1800 may include program instructions configured to implement some or all of the techniques described herein for in-stream rolling shutter compensation

(according to any of the embodiments described herein). For example, memory 1800 may include code to implement and/or execute any or all of the components of image processor 120.

5 [0089] In various embodiments, image processor 120, and/or individual sub-modules of image processor 120 may each be implemented in any of various programming languages. For example, in one embodiment, image processor 120 may be JAVA based, while in another embodiment, image processor 120 may be written using the C or C++ or assembly programming languages. Moreover, in some embodiments, different portions or sub-modules of image processor 120 may be implemented using different programming languages.

10 [0090] While various systems and methods have been described herein with reference to, and in the context of, specific embodiments, it will be understood that these embodiments are illustrative and that the scope of the disclosure is not limited to these specific embodiments. Many variations, modifications, additions, and improvements are possible. For example, the blocks and logic units identified in the description are for understanding the described embodiments and not meant to limit the disclosure. Functionality may be separated or combined  
15 in blocks differently in various realizations of the systems and methods described herein or described with different terminology.

[0091] These embodiments are meant to be illustrative and not limiting. Accordingly, plural instances may be provided for components described herein as a single instance. Boundaries between various components, operations and data stores are somewhat arbitrary, and particular  
20 operations are illustrated in the context of specific illustrative configurations. Other allocations of functionality are envisioned and may fall within the scope of claims that follow. Finally, structures and functionality presented as discrete components in the exemplary configurations may be implemented as a combined structure or component. These and other variations, modifications, additions, and improvements may fall within the scope of the disclosure as  
25 defined in the claims that follow.

[0092] Although the embodiments above have been described in detail, numerous variations and modifications will become apparent once the above disclosure is fully appreciated. It is intended that the following claims be interpreted to embrace all such variations and modifications.

30

**WHAT IS CLAIMED IS:**

1. A system, comprising:

5 a camera image sensor; and

an image processing unit configured to:

while receiving image data for an image from a camera image sensor:

10

accumulate a strip of the received image data comprising a plurality of pixel rows of the image;

15

determine, for each of a plurality of blocks across the strip of image data, a graphical profile for the given block;

20

determine, for each of the plurality of blocks, an offset for the given block based on comparing the determined graphical profiles with graphical profiles for a previous image;

25

determine an offset vector for the strip based on the determined offsets for the plurality of blocks; and

adjust the image data for the strip, according to the offset vector.

2. The system of claim 1, wherein to determine the graphical profile for a respective block, the image processing unit is further configured to:

30

sum graphical values for a plurality of pixels in each row of the block; and

sum graphical values for a plurality of pixels in each column of the block.

3. The system of claims 1 or 2, wherein the graphical values comprise luminance values for the plurality of pixels.

4. The system of claim 1, wherein to determine the graphical profile for a given block, the image processing unit is further configured to:

average graphical values for a plurality of pixels in each column of the block.

5. The system of any of claims 1 - 4, wherein to determine an offset for a given block,  
10 the image processing unit is further configured to determine one or both of:

a horizontal offset for the given block based on a comparison of a horizontal component of the graphical profile to graphical profiles for one or more other blocks of the previous image; and

15 a vertical offset for the given block based on a comparison of a vertical component of the graphical profile to graphical profiles for the one or more other blocks of the previous image.

20           6. The system of any of claims 1 - 5, wherein to determine an offset for a given block, the image processing unit is further configured to compare multiple possible matches between the graphical profile for the given block and the graphical profiles for the previous image.

7. The system of claim 6, wherein to compare among multiple possible matches, the  
25 image processing unit is further configured to compare, for each possible match, a sum of  
absolute differences between graphical profiles for a pixel location within the given block and  
graphical profiles for a location within one of the blocks of the previous image.

8. The system of claim 7, wherein to compare among multiple possible matches, the  
30 image processing unit is further configured to:

determine that a given possible match is a good fit if the sum of absolute differences between the pixel location with the given block and the pixel location with the one of the blocks of the previous image is smaller than a predetermined threshold; and

determine that a given possible match is a unique fit if the sum of absolute differences between the pixel location of the given block and the pixel location of the previous image is smaller, by a predetermined threshold, than respective sums of absolute differences for others of the multiple possible matches.

9. A method, comprising:

performing, by an image processor:

while receiving image data for an image from a camera image sensor:

accumulating a strip of the received image data comprising a plurality of pixel rows of the image;

for each of a plurality of blocks across a strip of image data for the image:

determining a graphical profile for the given block;

determining an offset for the given block based on comparing the determined graphical profiles with graphical profiles for a previous image;

determining an offset vector for the strip based on the determined offsets for the plurality of blocks; and

adjusting image data for the strip, according to the offset vector.

10. The method of claim 9, wherein said determining a graphical profile comprises:

summing graphical values for a plurality of pixels in each row of the block; and

summing graphical values for a plurality of pixels in each column of the block.

11. The method of claims 9 or 10, wherein the graphical profile for the given block comprises luminance data for the given block.

12. The method of any of claims 9 - 11, wherein said determining an offset vector for the  
5 strip comprises combining the respective offsets for one or more of the blocks of the horizontal strip.

13. The method of any of claims 9 - 12, wherein said adjusting comprises omitting one or  
more columns of pixels from the image data to compensate for a horizontal component of the  
10 offset vector.

14. The method of any of claims 9 - 13, wherein said adjusting comprises combining one  
or more rows of pixels in the image data to compensate for a vertical component of the offset  
vector.

15

15. The method of any of claims 9 - 13, wherein said warping comprises duplicating or  
resampling one or more rows of pixels in the image data to compensate for a vertical component  
of the offset vector.




FIG. 1




FIG. 2



**FIG. 3**  
SUBSTITUTE SHEET (RULE 26)




FIG. 4




FIG. 5




FIG. 6




FIG. 7





FIG. 8



# INTERNATIONAL SEARCH REPORT

International application No  
PCT/US2014/051792

A. CLASSIFICATION OF SUBJECT MATTER  
INV. H04N5/232  
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)  
G06T 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)

EP0-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         | US 2009/160957 A1 (DENG YINING [US] ET AL)<br>25 June 2009 (2009-06-25)                                                 | 1-4,6-15              |
| Y         | abstract<br>paragraphs [0001], [0005], [0023],<br>[0024], [0027], [0028], [0029],<br>[0030], [0053]<br>figure 2         | 5                     |
| Y         | -----<br>US 2006/140602 A1 (KURATA TOHRU [JP] ET<br>AL) 29 June 2006 (2006-06-29)<br>abstract<br>claim 6                | 5                     |
| A         | -----<br>US 2013/044230 A1 (ZHOU JIANPING [US])<br>21 February 2013 (2013-02-21)<br>the whole document<br>-----<br>-/-- | 1-15                  |

☒ 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

15 October 2014

Date of mailing of the international search report

23/10/2014

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

Lauri, Lauro

# INTERNATIONAL SEARCH REPORT

International application No

PCT/US2014/051792

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

| Category* | Citation of document, with indication, where appropriate, of the relevant passages                             | Relevant to claim No. |
|-----------|----------------------------------------------------------------------------------------------------------------|-----------------------|
| A         | US 2013/044241 A1 (ZHOU JIANPING [US])<br>21 February 2013 (2013-02-21)<br>the whole document                  | 1-15                  |
| A         | -----<br>US 2009/201383 A1 (SLAVIN KEITH R [US])<br>13 August 2009 (2009-08-13)<br>the whole document<br>----- | 1-15                  |

# INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2014/051792

| Patent document<br>cited in search report | Publication<br>date | Patent family<br>member(s)                           | Publication<br>date                    |
|-------------------------------------------|---------------------|------------------------------------------------------|----------------------------------------|
| US 2009160957 A1                          | 25-06-2009          | NONE                                                 |                                        |
| US 2006140602 A1                          | 29-06-2006          | JP 4389779 B2<br>JP 2006186481 A<br>US 2006140602 A1 | 24-12-2009<br>13-07-2006<br>29-06-2006 |
| US 2013044230 A1                          | 21-02-2013          | NONE                                                 |                                        |
| US 2013044241 A1                          | 21-02-2013          | NONE                                                 |                                        |
| US 2009201383 A1                          | 13-08-2009          | JP 4727735 B2<br>JP 2009189005 A<br>US 2009201383 A1 | 20-07-2011<br>20-08-2009<br>13-08-2009 |