



- (51) International Patent Classification:  
*H04N 5/232* (2006.01)
- (21) International Application Number:  
PCT/EP2012/060315
- (22) International Filing Date:  
31 May 2012 (31.05.2012)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:  
1109071.9 31 May 2011 (31.05.2011) GB
- (71) Applicant (for all designated States except US): SKYPE [IE/IE]; 70 Sir John Rogerson's Quay, Dublin 2 (IE).
- (72) Inventors; and
- (75) Inventors/Applicants (for US only): JEFREMOV, Andrei [EE/SE]; Veckovagen 59-274, S-177 62 Jarfalla (SE). ZHAO, David [SE/SE]; Honnorsgatan 24, S-170 69 Solna

(SE). CARLSSON, Pontus [SE/SE]; Byggmastarvagen 3, S-168 32 Bromma (SE). SABLIN, Sergey [RU/SE]; Svartviksslingan, S-167 38 Bromma (SE).

(74) Agents: DRIVER, Virginia Rozanne et al.; Page White & Farrer, Bedford House, John Street, London, Greater London WC1N 2BF (GB).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, 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, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, 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,

[Continued on next page]

(54) Title: VIDEO STABILISATION

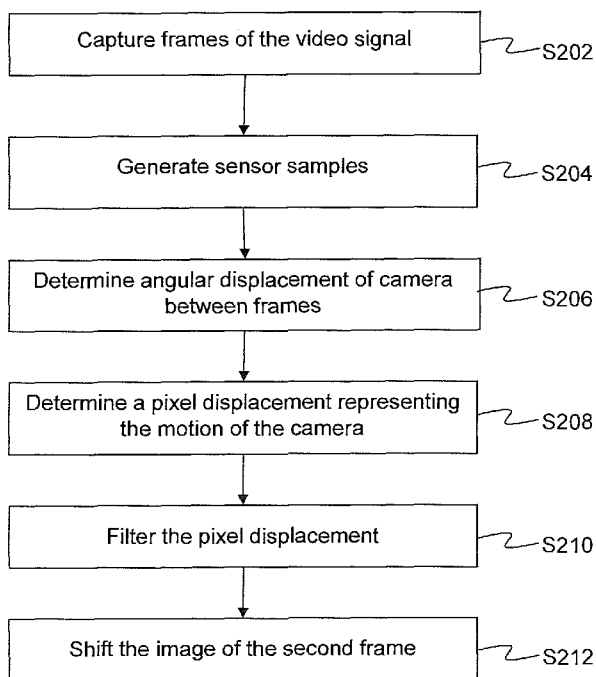


Fig 2

(57) Abstract: Method, device and computer program product for stabilising a video signal. A plurality of frames of the video signal are captured using a camera. A motion sensor associated with the camera is used to generate a plurality of samples representing motion of the camera. The samples are used to determine a displacement of the camera between a first time and a second time, wherein the first time corresponds to an exposure time midpoint of a first frame of the video signal and the second time corresponds to an exposure time midpoint of a second frame of the video signal. The determined displacement is used to compensate for motion in the video signal between the first and second frames caused by the motion of the camera, to thereby stabilise the video signal.



GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, 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, ML, MR, NE, SN, TD, TG).

**Published:**

— *without international search report and to be republished upon receipt of that report (Rule 48.2(g))*

## VIDEO STABILISATION

### Field of the Invention

- 5 The present invention relates to stabilisation of a video signal. In particular, the present invention relates to capturing frames of a video signal using a camera and compensating for motion of the camera to thereby stabilise the video signal.

### Background

10

Cameras can be used to capture a sequence of images to be used as frames of a video signal. Cameras may be fixed to stable objects, for example a camera may be mounted on a stand such as a tripod to thereby keep the camera still while the video frames are captured. However, often cameras may be implemented in mobile devices and are not necessarily mounted to fixed objects, for example a camera may be held, or may be on a moving object such as a vehicle. Movement of the camera while the camera is capturing frames of a video signal may result in unwanted movement in the video signal itself.

15

- 20 Image stabilisation is a method that can be used to compensate for the unwanted movement in a video signal. Some systems perform motion estimation in order generate motion vectors for use by an image stabilisation process. One such system is described in "Online Video Stabilization Based on Particle Filters" by Junlan Yang et. al. Image stabilisation algorithms may consist of three main parts: motion estimation, motion smoothing and motion compensation. A motion estimation block may estimate local motion vectors within the video signal and on the basis of these local estimates calculate a global motion vector. A motion smoothing block may then deal with filtering of the estimated global motion vector in order to smooth the calculated value and prevent large and undesirable differences between motion vectors calculated previously. A motion compensation block may then shift an image in the opposite direction to the filtered global motion vector to thereby stabilise the video signal. The motion compensation block may take into account sophisticated transformations like rotation, warping or zooming.
- 25
- 30

It can require large amounts of processing resources to perform image stabilisation based on motion vectors as described above. This can be a problem when the video signal is to be stabilised in real time, i.e. when a stabilised version of the video signal is to be used (e.g. transmitted in a video call or output from a device) at the same time as it is being captured by the camera. This can also be a problem when the device which is performing the image stabilisation is a small, mobile device such as a mobile telephone in which the processing resources are limited.

10

In recent years, motion sensors have become simpler and cheaper to manufacture and the size of motion sensors has reduced significantly. It is now feasible to implement motion sensors in mobile devices. Motion sensors generate samples representing the motion of the sensor. Two prior art documents: "Using Sensors for Efficient Video Coding in Hand-held devices" by Andy L. Lin and "Accelerometer Based Digital Video Stabilization for General Security Surveillance Systems" by Martin Drahanský et. al, mention the possibility of using data from motion sensors for stabilization of a video signal.

15

## 20 Summary

The inventors have realised that the implementation of using data from motion sensors for stabilization of a video signal can be improved compared to the systems of the prior art. In particular, the inventors have considered factors such as: handling non-uniformly sampled sensor data; synchronization of sampled sensor data with the timing of frames of the video signal taking the shutter speed of the camera into account; fast pixel shifting of an image of one of the frames of the video signal by pointer modification; and how to select sensor sampling rate. These factors have not been considered in the prior art cited in the background section above.

25

30

According to a first aspect of the invention there is provided a method of stabilising a video signal, the method comprising: capturing a plurality of frames of the video signal using a camera; using a motion sensor associated with the

camera to generate a plurality of samples representing motion of the camera; using the samples to determine a displacement of the camera between a first time and a second time, wherein the first time corresponds to an exposure time midpoint of a first frame of the video signal and the second time corresponds to an exposure time midpoint of a second frame of the video signal; and using the determined displacement to compensate for motion in the video signal between the first and second frames caused by the motion of the camera, to thereby stabilise the video signal.

5

10 Time stamps provided by the camera for the frames of the video signal may relate to the time at which a shutter of the camera closes (i.e. the end time of the frame). However, the inventors have identified that in order to effectively determine the displacement (e.g. an angular displacement) between the first and second frames of the video signal using data from the motion sensor (e.g. a rotational motion sensor), it is advantageous to determine the displacement of the camera between the midpoints of the exposure times of the frames. Using the midpoints of the exposure times of the frames provides a more accurate representation of the displacement of the camera between the first and second frames. When mapping camera displacement to pixel displacement, this method provides better estimation of pixel displacement than using the end times of the frames.

15

20

In preferred embodiments, the motion of the camera is rotational motion, the motion sensor is a rotational motion sensor and the displacement of the camera is an angular displacement of the camera. The using of the samples to determine an angular displacement of the camera between a first time and a second time may comprises: determining an angular velocity of the camera using the samples; and integrating the determined angular velocity over time from the first time to the second time to thereby determine the angular displacement of the camera between the first time and the second time.

25

30

In one embodiment, the angular displacement of the camera between the first time and the second time is given by the equation:

$$\Delta\theta = \int_{t_1 - 0.5e_1}^{t_2 - 0.5e_2} \omega(t) dt$$

where  $\Delta\theta$  is the angular displacement of the camera,  $t_1$  is end time of the first frame,  $t_2$  is end time of the second frame,  $e_1$  is the exposure time of the first frame,  $e_2$  is the exposure time of the second frame and  $\omega(t)$  is the angular velocity of the camera determined using the samples.

5

The integrating of the determined angular velocity may comprise interpolating the angular velocity between the times at which the samples are generated and the method may further comprise delaying the video signal thereby allowing the angular velocity of the camera at the second time to be determined by said  
10 interpolating the angular velocity between the times at which the samples are generated.

The method may further comprise extrapolating the angular velocity determined using the samples to thereby determine the angular velocity of the camera at  
15 the second time.

The sample rate of the samples generated using the motion sensor may be higher than the frame rate of the video signal. Furthermore, the camera and the motion sensor may be situated within a mobile device.

20

The using of the determined displacement to compensate for motion in the video signal between the first and second frames caused by the motion of the camera may comprise: determining a pixel displacement representing motion in the video signal between the first and second frames caused by the determined  
25 displacement of the camera; filtering the pixel displacement; and shifting the image of at least one of the first and second frames in accordance with the filtered pixel displacement to thereby stabilise the video signal. The filtering of the pixel displacement may comprise: determining an accumulated pixel displacement based on said determined pixel displacement for the second  
30 frame; and determining a filtered accumulated pixel displacement for the second frame based on a weighted sum of the determined accumulated pixel

displacement for the second frame and a filtered accumulated pixel displacement for the first frame.

5 The method may further comprise adding a time offset to at least one of (i) the captured plurality of frames, and (ii) the generated plurality of samples, such that the timing of the captured plurality of frames matches the timing of the generated plurality of samples.

10 According to a second aspect of the invention there is provided a device for stabilising a video signal, the device comprising: a camera configured to capture a plurality of frames of the video signal; a motion sensor, associated with the camera, configured to generate a plurality of samples representing motion of the camera; a displacement determining block configured to use the samples to determine a displacement of the camera between a first time and a second time, wherein the first time corresponds to an exposure time midpoint of a first frame of the video signal and the second time corresponds to an exposure time midpoint of a second frame of the video signal; and a motion compensation block configured to use the determined displacement to compensate for motion in the video signal between the first and second frames caused by the motion of the camera, to thereby stabilise the video signal.

The motion sensor may be a gyroscopic motion sensor. The device may be a mobile device.

25 According to a third aspect of the invention there is provided a method of stabilising a video signal, the method comprising: capturing a plurality of frames of the video signal using a camera; determining a pixel displacement representing motion in the video signal between first and second frames of the video signal caused by motion of the camera; and shifting an image of at least one of the first and second frames in accordance with the pixel displacement to thereby stabilise the video signal, wherein the image comprises a stride value indicating memory space of the image, a plurality of pixel values, a pointer indicating the position of a first pixel of the image, and a width value indicating

the width of the image, wherein said shifting of the image comprises adjusting the pointer and the width value without adjusting the stride value.

5 In this way, the image may be shifted and resized by simply changing the pointer and the width value whilst keeping the stride value constant. In this way, no copying of data in memory is necessary. In other words, instead of copying a crop area into a new memory area (which can be a complex process), an image representation is used that allows for independent width and stride values. In this way, a new image can be created by changing the pointer  
10 and the width value while the stride is kept intact (which is a simpler process than copying a crop area into a new memory area).

The image may comprise a plurality of image planes which are each represented by a respective plurality of pixel values, a respective pointer  
15 indicating the position of a first pixel of the image plane, and a respective width value indicating the width of the image plane, wherein the image planes may be shifted together by adjusting their respective pointers and width values without adjusting the stride value.

20 The method may further comprise rounding the determined pixel displacement to an integer-pixel displacement.

The method may further comprise: using a motion sensor associated with the camera to generate a plurality of samples representing motion of the camera;  
25 and using the samples to determine a displacement of the camera between a first frame of the video signal and a second frame of the video signal, wherein the determined displacement is used to determine said pixel displacement.

The method may further comprise filtering the pixel displacement.  
30

According to a fourth aspect of the invention there is provided a device for stabilising a video signal, the device comprising: a camera configured to capture a plurality of frames of the video signal; a pixel displacement determining block configured to determine a pixel displacement representing motion in the video

signal between first and second frames of the video signal caused by motion of the camera; and an image shifting block configured to shift an image of at least one of the first and second frames in accordance with the pixel displacement to thereby stabilise the video signal, wherein the image comprises a stride value  
5 indicating memory space of the image, a plurality of pixel values, a pointer indicating the position of a first pixel of the image, and a width value indicating the width of the image, wherein the image shifting block is configured to shift the image by adjusting the pointer and the width value without adjusting the stride value.

10

According to a fifth aspect of the invention there is provided a method of stabilising a video signal, the method comprising: capturing a plurality of frames of the video signal using a camera; determining a portion of motion of the camera occurring whilst a shutter of the camera is closed and not whilst the  
15 shutter of the camera is open; determining a pixel displacement representing motion in the video signal between first and second frames of the video signal caused by, and corresponding to, the portion of motion of the camera; and shifting an image of at least one of the first and second frames in accordance with the pixel displacement to thereby stabilise the video signal.

20

In this way, image stabilisation may be reduced when motion blur is present in the video signal. For high levels of camera motion, motion blur will be present in the frames of the video signal for devices without physical stabilisation, such as optical or mechanical stabilisation. Motion blur is caused when the camera  
25 moves whilst the shutter of the camera is open and capturing one frame of the video signal. If image stabilisation were applied for frames containing motion blur, the resulting video signal will contain motion blur but no motion which can be detrimental because a user viewing the video signal may perceive this to look strange and/or unnatural. The inventors have realised that it may be  
30 beneficial to have no compensation for motion that is made during the time interval when the camera shutter is open. If the exposure time (i.e. the time for which the shutter is open) is very short, then having no compensation for motion that occurs whilst the shutter is open will not make a significant difference. However, if the exposure time (i.e. the time for which the shutter is open) is very

long, then little stabilization will be applied to the video signal. Furthermore, motion blur will be accompanied by the corresponding motion in the video signal, which may be perceived as more natural to a user viewing the video signal.

5

The method may further comprise: using a motion sensor associated with the camera to generate a plurality of samples representing the portion of motion of the camera; and using the samples to determine a displacement of the camera between the first frame of the video signal and the second frame of the video signal, wherein the determined displacement is used to determine said pixel displacement.

According to a sixth aspect of the invention there is provided a device for stabilising a video signal, the device comprising: a camera configured to capture a plurality of frames of the video signal; a motion determining block configured to determine a portion of motion of the camera occurring whilst a shutter of the camera is closed and not whilst the shutter of the camera is open; a pixel displacement determining block configured to determine a pixel displacement representing motion in the video signal between first and second frames of the video signal caused by, and corresponding to, the portion of motion of the camera; and an image shifting block configured to shift an image of at least one of the first and second frames in accordance with the pixel displacement to thereby stabilise the video signal.

According to a seventh aspect of the invention there is provided a computer program product for stabilising a video signal, the computer program product being embodied on a non-transient computer-readable medium and configured so as when executed on a processor of a device to perform the operations of any of the methods described herein.

30

There are described herein methods for using motion sensor data to remove camera shaking (and other motion of the camera) from a video signal to thereby stabilize the video signal, for example for use in a mobile device. The methods may comprise synchronising sensor data with camera frames and, in particular,

determining the angular displacement at the midpoint of the exposure time. The methods may comprise filtering image displacement values using an exponential filter suitable for use with non-uniformly sampled sample data from the motion sensor. The methods may comprise modifying the position of a pointer for shifting and/or cropping an image of a frame of the video signal based on determined motion of the camera. The methods may comprise adapting image stabilisation for frames of the video signal containing motion blur.

## 10 Brief Description of the Drawings

For a better understanding of the present invention and to show how the same may be put into effect, reference will now be made, by way of example, to the following drawings in which:

15

Figure 1 shows a device according to a preferred embodiment;

Figure 2 is a flow chart for a process of stabilising a video signal according to a preferred embodiment;

Figure 3 is an exemplary graphical representation of a shutter position of a camera over time;

Figure 4 is an exemplary graphical representation of an angular velocity of a camera over time; and

Figure 5 is a representation of an image before and after a cropping process.

## 25 Detailed Description of Preferred Embodiments

Preferred embodiments of the invention will now be described by way of example only.

30 Figure 1 shows a device 100 according to a preferred embodiment. The device 100 may for example be a mobile device such as a mobile telephone or other hand-held device. The device 100 comprises a camera 102, a motion sensor 104, a CPU 106 and a memory 108. The camera 102 is configured to capture images. The captured images may be used to form a video signal, whereby

each image is used as a frame of the video signal and the images are captured at the frame rate of the video signal. The frame rate may be, for example, 25 frames per second although the camera may operate at a different frame rate. The minimum frame rate to achieve the perception of a moving image is about  
5 fifteen frames per second, but this may depend upon the person viewing the video signal and upon the content that is in the video signal (i.e. how much motion there is in the subject matter of the video signal). The motion sensor 104 is configured to generate samples representing the motion of the device 100. Since both the motion sensor 104 and the camera 102 are in the device  
10 100, they are associated with each other such that the samples generated by the motion sensor 104 can be used to represent the motion of the camera 102. The CPU 106 is configured to perform computational processes on the device 100 as is known in the art. The memory 108 is used to store data in the device 100 as is known in the art. The blocks 102, 104, 106 and 108 can communicate  
15 with each other by sending data via a bus of the device 100 (not shown in Figure 1) as is known in the art.

With reference to Figure 2 there is now described a process of stabilising a video signal using the device 100 according to a preferred embodiment. In step  
20 S202 the camera 102 captures images to be used as frames of a video signal. For example, the camera 102 may have an array of light sensors which record the level of light that is incident on the sensors during the time allocated to a frame of the video signal. A shutter of the camera 102 is used to separate the frames in time, such that during each frame the shutter is open for a period of  
25 time and closed for another period of time. The captured frames of video signal are provided to a pre-processor (e.g. implemented in a processing block by the CPU 106). The pre-processor operates to stabilise the images in the frames of the video signal before the frames are encoded using a video encoding  
30 technique as is known in the art.

In step S204, while the camera 102 is capturing frames of the video signal, the motion sensor 104 generates samples representing the motion of the device 100. For example, the motion sensor 104 may be a rotational motion sensor such as a gyroscope. The gyroscope 104 measures angular velocity of the

device 100 and outputs samples representing the angular velocity at particular intervals. The intervals may, or may not, be regular intervals. Preferably, but not necessarily, on average the sample rate of the samples output from the gyroscope 104 is higher than the frame rate of the video signal. For example, the sample rate output from the gyroscope 104 may be 60 samples per second, which reflects the maximum usual shaking frequency of the device 100 and is currently independent of frame rate. The samples generated by the gyroscope 104 are provided to the pre-processor.

10 In step S206 the angular displacement of the camera 102 between two frames (frame 1 and frame 2) of the video signal is determined. This determination may be performed by a processing block of the CPU 106. The inventors have identified that in order to effectively determine the angular displacement between the two frames using data from the gyroscope 104, it is useful to integrate the angular velocity over the time interval between the midpoints of the exposure times of the frames captured by the camera 102. The inventors have also determined that this can be particularly problematic as it may not be possible to synchronise the sampling rate of the gyroscope 104 with the frame rate of the camera 102, particularly when:

- 20 - the camera 102 is arranged to adjust the exposure times in dependence on the available light (which many cameras are);
- the time stamps for the frames of the video signal provided by the camera 102 relate to the times at which the shutter closes (i.e. the end times of the frames, as opposed to the midpoints of the exposure times of the frames); and
- 25 - the gyroscope data is not available at the midpoint of the exposure time of the frames.

As described above, the pre-processor receives video frames from the camera 102 and also receives the samples from the gyroscope 104. The samples from the gyroscope 104 are provided to the pre-processor (e.g. at regular intervals), at a rate at least equivalent to the frame rate of the video signal captured by the camera 102. Using a higher sampling rate in the gyroscope 104 gives more accurate angle estimates but can be more costly in terms of CPU usage.

A time stamp,  $t_1$ , provided by the camera 102 to a first frame (frame 1) of the video signal indicates the end time of the frame, i.e. the time at which the shutter of the camera 102 is closed to end frame 1. Similarly, a time stamp,  $t_2$ , provided by the camera 102 to a second frame (frame 2) of the video signal indicates the end time of the frame, i.e. the time at which the shutter of the camera 102 is closed to end frame 2. In order to determine the angular displacement ( $\Delta\theta$ ) of the device 100 between the first frame and the second frame, rather than using the time stamps of the frames to denote the times of the frames, it is more accurate to use the midpoints of the exposure time of frame 1 and frame 2. The exposure times of the first and second frames are denoted by  $e_1$  and  $e_2$ . The angular displacement is determined by integrating the angular velocity (represented by the samples output from the gyroscope 104) of the device 100 between a time  $t_1 - 0.5e_1$  and a time  $t_2 - 0.5e_2$ . Therefore the angular displacement between frame 1 and frame 2 is given by:

$$\Delta\theta = \theta(t_2) - \theta(t_1) = \int_{t_1 - 0.5e_1}^{t_2 - 0.5e_2} \omega(t) dt$$

Figure 3 is an exemplary graphical representation of a shutter position of the camera 102 over time. The shutter of the camera 102 closes at time  $t_1$  at the end of frame 1. The shutter re-opens again for the camera 102 to capture frame 2 and then closes at time  $t_2$  at the end of frame 2. The exposure time of frame 1 is shown as  $e_1$  and the exposure time of frame 2 is shown as  $e_2$  in Figure 3. The time over which the angular velocity is integrated is denoted  $T_{12}$  in Figure 3. It can be appreciated from looking at Figure 3 that integrating over the time  $T_{12}$  corresponds to integrating between the midpoint of the exposure time of the first frame (at time  $t_1 - 0.5e_1$ ) and the midpoint of the exposure time of the second frame (at time  $t_2 - 0.5e_2$ ). Figure 3 shows the open time of the shutter to be equal to the closed time of the shutter, but this is just one example. In some embodiments (implementing short exposure times) the time for which the shutter is open is shorter than the time for which the shutter is closed. In contrast, in other embodiments (implementing long exposure times) the time for which the shutter is open is longer than the time for which the shutter is closed.

Since the samples of the gyroscope 104 are not synchronised with the timings of the frames of the video signal captured by the camera 102, it might be the case that the gyroscope 104 does not generate samples at the midpoints of the frames (frame 1 and frame 2). In which case, the angular velocity of the device  
5 100 at the midpoints of the frames can be determined by interpolating the angular velocity represented by the samples generated by the gyroscope 104. The angular velocity is evaluated by interpolation at any time instant, and the midpoints of the exposure times of the frames define the integral interval used when calculating the angular displacement according to the equation above.

10

Figure 4 is an exemplary graphical representation of an angular velocity of the camera 102 over time. The samples representing the angular velocity of the device 100 generated by the gyroscope 104 are shown in Figure 4 as samples 402, 404, 406, 408 and 410. It can be seen that in the example shown in Figure  
15 4 the timings of the samples of the gyroscope 104 are not regular. For example, the time between the samples 404 and 406 is shorter than the time between the samples 406 and 408. The dotted line connecting the samples in Figure 4 shows the value of the angular velocity that can be determined as a function of time by interpolating the angular velocity represented by the samples  
20 generated by the gyroscope 104. The interpolated angular velocity (shown by the dotted line) can be integrated between times  $(t_1 - 0.5e_1)$  and  $(t_2 - 0.5e_2)$  in order to determine the angular displacement of the camera 102 between the first and second frames. Figure 4 shows a simple linear interpolation between the samples from the gyroscope 104. In other embodiment, more advanced  
25 interpolation could be used.

There may arise a situation in which a frame to be stabilised is received at the pre-processor after the latest sample from the gyroscope 104. For example, when the frame 2 is captured at the camera 102 the frame 2 may be received at  
30 the pre-processor before any samples from the gyroscope have been generated subsequent to the midpoint of the exposure time of the frame 2  $(t_2 - 0.5e_2)$ . For example frame 2 may be received at the pre-processor before the sample 410 shown in Figure 4. In this situation, delay may be introduced into the video stream, in order for the sample 410 to be received at the pre-processor before

the frame 2 is processed, thereby allowing the angular velocity at time  $(t_2-0.5e_2)$  to be determined before the frame 2 is processed by the pre-processor. Alternatively, the angular velocity may be extrapolated from the previously received samples from the gyroscope 104 in order to determine the angular velocity of the device 100 at the time  $(t_2-0.5e_2)$ .

In the case of no motion of the camera 102 (e.g. for fixed placement of the device 100), the gyroscope 104 may be disabled in order to save battery life. The state of no motion can be determined by feedback from a video encoder which encodes the video signal subsequent to the image stabilisation method described herein and implemented by the pre-processor. The video encoder may perform motion estimation as part of the encoding process and as such can determine whether the camera is moving. A state of motion can also be determined and used to enable the gyroscope 104 when the camera 102 is moved. When the device 100 operates in the state of no motion, the motion sensor 104 may be polled at a slow interval to determine whether the device 100 has started moving again. There may be computationally cheaper ways to determine when the device 100 starts moving, depending on hardware and Application Programming Interfaces (APIs) implemented in the operating system of the device 100.

The timings of the operation of the hardware used for the camera 102 and for the gyroscope 104 might not match. This may be because the camera 102 and the gyroscope 104 are implemented in independent hardware chips. Therefore it may be beneficial to add an offset to the time stamps of either (or both) the samples generated by the gyroscope 104 and the frames of the video signal. In this way the timing of the samples from the gyroscope 104 can be matched with the timing of the frames of the video signal correctly. The offset may be constant for a particular combination of hardware chips. Therefore a delay may be computed offline and used at the device 100 without incurring a processing penalty for the method described herein.

In step S208 a pixel displacement representing the motion of the camera 102 is determined. In general, a rotation of the camera 102 results in an

approximately constant pixel displacement across the image of a frame of the video signal, independent of distances to objects in the image. This is in contrast to linear camera motion, for which pixel displacement is a function of the distance to the object. A function (or algorithm) mapping the rotation of the device 100 to a pixel displacement depends on parameters of the camera 102 (e.g. focal length and width of lens of the camera 102) and the resolution of the images captured by the camera 102. Encoder feedback can be useful to determine the accuracy of the samples generated by the gyroscope 104 and to adapt the mapping algorithm. There are also some cases of motion and object placement where the stabilization model described herein based on the samples from the gyroscope 104 is not accurate (e.g. for rotation of the camera 102 around a user's face, the user's face may be stable in the middle of the frame but the gyroscope 104 detects rotation and therefore the stabilisation process will attempt to stabilise the background) which may be detected by the encoder and fed back to the stabilisation algorithm. In this way the stabilisation algorithm can be adapted.

The pixel displacement determined in step S208 represents the motion in the images of the frames of the video signal resulting from the motion of the camera 102 (as opposed to motion in the subject matter of the images). In this way, the pixel displacement determined in step S208 represents unwanted motion in the images of the frames of the video signal.

In step S210 the pixel displacement determined in step S208 is filtered. This is done in order to smooth the changes that are applied to the video signal in the image stabilisation process over time to thereby provide a smoother stabilised video signal. The filter used to filter the pixel displacement can be designed in different ways, depending on, for example, the resolution of the images captured by the camera 102, the acceptable delay which may be applied to the video signal, and the allowed amount of cropping which can be applied to the images of the original video signal received at the pre-processor from the camera 102. For example, higher resolution video frames may benefit from a larger filter attenuation of high frequency changes to the pixel displacement

applied in the image stabilisation process. On the other hand, the amount of cropping sets a hard limit to the maximum filter attenuation.

5 An exponential filter may be used which filters the pixel displacements according to the equation:

$$x\_filt(n) = (1-w)*x\_filt(n-1) + w*x(n),$$

where n represents the frame number of the video signal, x represents the accumulated displacement (or "position"), according to the pixel displacement determined in step S208, and x\_filt represents the filtered accumulated displacement which is subsequently used to determine how to align the input image in order to stabilize it as described in more detail below. In this way the filter acts as an exponential filter. When motion stops, x\_filt – x will converge to zero which implies no shifting of the image. The filter smoothes out changes to the determined pixel displacement over time, by basing the filtered pixel displacements on the corresponding filtered pixel displacement of the previous frame as well as on the pixel displacement determined for the current frame in step S208. The weighting applied to the filtered pixel displacement of the previous frame is (1-w) whereas the weighting applied to the pixel displacement determined for the current frame is w. Therefore adjusting the weighting parameter, w, will adjust how responsive the filter is to changes in the pixel displacement (x). A recursive (Infinite Impulse Response (IIR)) filter is more suited than a Finite Impulse Response (FIR) filter when the output x\_filt is clipped to be in the range [x - crop, x + crop] as the clipped value is fed back to the filter loop and makes subsequent output of x\_filt less prone to clipping.

25

The weighting parameter, w, is adapted to the resolution and instant frame rate of the video signal to obtain a constant physical cut-off frequency, which is measured in Hertz. If the filter were an ideal filter then the physical cut-off frequency would define the highest frequency component of changes to x which will be incorporated into x\_filt. Changes to x which have higher frequency than the cut-off frequency will be attenuated by an ideal filter and will not be present in x\_filt. However, the filter is not an ideal filter and as such the cut-off frequency defines the highest frequency for which the attenuation applied by the filter is below 3dB. So for non-ideal filters there will be some attenuation below

30

the cut-off frequency and there will not be perfect attenuation above the cut-off frequency. The filter output is clipped so that that the difference between  $x_{\text{filt}}$  and  $x$  is not larger than the frame cropping size.  $w$  is adapted so that the physical cut-off frequency is constant, e.g. 0.5 Hz. From the filter transfer function, a function  $w(f_c, f_s)$  can be derived that maps a physical cut-off frequency  $f_c$  to  $w$ . When the sampling frequency (frame rate)  $f_s$  changes,  $w$  also changes even though  $f_c$  is constant. The filter according to the filter equation above is well suited for instant changing of the cut-off frequency (changing  $w$ ), compared to other filters.

10

In step S212 the image of the second frame (frame 2) is shifted using the filtered pixel displacement from step S210. In this way the motion in the image of the second frame (relative to the first frame) due to the motion of the camera 102 is attenuated. In other words, the filtered pixel displacement is used to compensate for the motion in the video signal between the first and second frames caused by the motion of the camera, to thereby stabilise the video signal.

The filtered pixel displacements are rounded to full-pixel displacements (i.e. integer-pixel displacements). This allows a simple method to be employed to shift the image of the second frame. The image is represented using a stride value indicating memory space of the image, a plurality of pixel values, a pointer indicating the position of a first pixel of the image, and a width value indicating the width of the image. The shifting of the image comprises adjusting the pointer and the width value without adjusting the stride value. It can be seen that the width value is independent of the stride value which allows the width of the image to be changed without affecting the stride of the image. Therefore the memory space of the image (e.g. in the memory 108) does not need to be changed when the image is shifted (and/or resized). This means that no copying of data in the memory 108 is necessary with this approach. This is in contrast to a conventional method of cropping an image in which the crop area of the image is copied into a new memory area. Copying the crop area may be computationally complex which may be detrimental, particularly when the method is to be implemented on a mobile device in which the processing

resources available to the CPU 106 may be limited. With the method described herein, since the width value is independent of the stride value the new, shifted image can be created by changing the pointer and the width while the stride is kept intact.

5

The image may be represented by multiple image planes, for example a luma plane (Y) and two chroma planes (U and V). The image planes of the input image may be shifted and resized by simply changing the pointers to the luma and chroma planes, thereby modifying the width of the image planes whilst  
10 keeping the stride intact. The image planes are shifted by the same amount to ensure that the shifted image planes can be used together to represent the shifted image.

In order for this image shifting process to be implemented, the image planes  
15 require respective pointers, i.e. they cannot all be represented by the same, single pointer. Furthermore, as described above, it is necessary that the image has independent width and stride values.

Figure 5 is a representation of an image before and after a shifting and cropping  
20 process. The original image is denoted 502 and the shifted and cropped image is denoted 504. It can be seen that the stride value of the image is left unchanged, whereas the width of the image is reduced. Furthermore, the original pointer points to the top left pixel of the original image whereas the adjusted pointer points to the top left pixel of the shifted and cropped image  
25 (which is in a different position to the top left pixel of the original image). In this way the image is shifted and cropped simply by changing the width value and the pointer.

In summary of the method described above, e.g. with reference to Figure 2, the  
30 following stages are implemented in the pre-processor to stabilise the images of the frames of the video signal before the video signal is encoded with a video encoder:

1. the angular displacement of the camera 102 between frame 1 and frame 2 is estimated (step S206);

2. the estimated angular displacement is mapped to a pixel displacement of the image of frame 2 (step S208);

3. unintended motion in the image of frame 2 is removed by applying a filter to the sequence of pixel displacements (or to the accumulated pixel displacements as described above) (step S210); and

4. a stabilized image for frame 2 is created by shifting the image to the position calculated by the filter (step S212). The frame dimensions of the stabilized image for frame 2 are equal or less than the corresponding dimensions of the original image for frame 2. In other words, the stabilized images of the video signal are constructed by cutting out a moving border within the original images of the video signal captured by the camera 102.

In some embodiments, the image stabilization may be reduced when motion blur is present in the video signal. When high levels of motion are experienced by the camera 102, motion blur will be present in the video. Therefore, if the image stabilisation described herein is applied to the frames of the video signal, then the resulting video signal will contain motion blur but no motion, which a user may perceive as looking unnatural or weird.

In one embodiment no compensation for motion of the camera 102 is made during the time interval when the camera shutter is open (i.e. the exposure times of the frames of the video signal) e.g. in the time interval  $[t_1 - e_1, t_1]$ . If the exposure time is very short (i.e. the shutter is open for a much shorter time than it is closed for) then this will not make a significant difference. However, if the exposure time is very long (i.e. the shutter is open for a much longer time than it is closed for), then little image stabilization will be applied to the video signal. In this way it is ensured that motion blur will always have the corresponding motion in the video signal, which is perceived as looking more natural to a viewer of the video signal.

30

Determining the motion of the camera 102 from the samples generated by the gyroscope 104 and determining the shutter speed is particularly advantageous since this information can be used to determine whether or not to apply the

image stabilisation, in view of the motion blur that the video signal may experience, without adding extra computational complexity.

5 It is also possible to address the issue of motion blur with no corresponding motion using a pure software stabiliser, rather than the image stabiliser described above which uses the samples from the gyroscope 104 to determine the motion of the camera 102. Software motion estimation typically finds the motion vectors representing centres of shutter speed intervals. Analyzing motion blur and finding what kind of motion produced this blur is more complex  
10 with a software stabiliser implementation than when using the samples from the gyroscope 104 to determine the motion of the camera 102, in combination with shutter speed, as described above.

In the embodiments described above, the motion sensor 104 is a gyroscope  
15 which generates samples representing the rotational motion of the device 100. In other embodiments, the motion sensor 104 may sense other types of motion, such as translational motion and generate samples representing the translational motion of the device 100. These samples can be used in the same way as described above in relation to the rotational motion to stabilise the video  
20 signal. However, as described above, with translational motion the pixel displacement will depend on the distance to the object in the image and so this must be taken into account when determining the pixel displacements. For example, multiple accelerometers may be able to estimate rotational motion, and in this case accelerometers can be used without further modification. For  
25 more general translational stabilization, it may become more difficult to implement the method described herein since different areas in the image move by different amounts of pixels. However, if the distance to the object is constant (and known) it may be simple to implement the method with translation motion. Even where the distance to the objects is not constant (but is still known) it  
30 would be possible to implement the method with translation motion but extra complication is added in determining the pixel displacements caused by the translation motion of the camera 102.

- After stabilising the video signal the video signal is encoded using a video encoding process. The encoded video signal may be transmitted, e.g. as part of a video call to another user or as a broadcast signal. Therefore, it is important for the video signal to be able to be stabilised and encoded in real-time (i.e. with very little delay) for use in events such as video calls, or other communication events where users are perceptually very aware of delay in the signals. Alternatively, the encoded video signal could be stored at the device 100, e.g. in the memory 108.
- 5
- 10 The method steps S206, S208, S210 and S212 could be implemented at the device 100 in software or in hardware. For example, the CPU 106 may execute processing blocks to implement the steps S206, S208, S210 and S212. For example, a computer program product for stabilising a video signal may be provided, which can be stored in the memory 108 and executed by the CPU
- 15 106. The computer program product may be configured so as when executed on the CPU 106 to perform the method steps S206, S208, S210 and S212. Alternatively, hardware blocks may be implemented in the device 100 to implement the steps S206, S208, S210 and S212.
- 20 Furthermore, while this invention has been particularly shown and described with reference to preferred embodiments, it will be understood to those skilled in the art that various changes in form and detail may be made without departing from the scope of the invention as defined by the appendant claims.

## Claims

1. A method of stabilising a video signal, the method comprising:  
 capturing a plurality of frames of the video signal using a camera;  
 5 using a motion sensor associated with the camera to generate a plurality of samples representing motion of the camera;  
 using the samples to determine a displacement of the camera between a first time and a second time, wherein the first time corresponds to an exposure time midpoint of a first frame of the video signal and the second time  
 10 corresponds to an exposure time midpoint of a second frame of the video signal; and  
 using the determined displacement to compensate for motion in the video signal between the first and second frames caused by the motion of the camera, to thereby stabilise the video signal.
- 15
2. The method of claim 1 wherein the motion of the camera is rotational motion, the motion sensor is a rotational motion sensor and the displacement of the camera is an angular displacement of the camera.
- 20
3. The method of claim 2 wherein said using the samples to determine an angular displacement of the camera between a first time and a second time comprises:  
 determining an angular velocity of the camera using the samples; and  
 integrating the determined angular velocity over time from the first time  
 25 to the second time to thereby determine the angular displacement of the camera between the first time and the second time.
4. The method of claim 3 wherein the angular displacement of the camera between the first time and the second time is given by the equation:

$$\Delta\theta = \int_{t_1 - 0.5e_1}^{t_2 - 0.5e_2} \omega(t) dt$$

- 30 where  $\Delta\theta$  is the angular displacement of the camera,  $t_1$  is end time of the first frame,  $t_2$  is end time of the second frame,  $e_1$  is the exposure time of the first

frame,  $e_2$  is the exposure time of the second frame and  $\omega(t)$  is the angular velocity of the camera determined using the samples.

5. The method of claim 3 or 4 wherein said integrating the determined  
5 angular velocity comprises interpolating the angular velocity between the  
times at which the samples are generated.

6. The method of claim 5 further comprising delaying the video signal  
10 thereby allowing the angular velocity of the camera at the second time to be  
determined by said interpolating the angular velocity between the times at  
which the samples are generated.

7. The method of any of claims 3 to 5 further comprising extrapolating the  
15 angular velocity determined using the samples to thereby determine the  
angular velocity of the camera at the second time.

8. The method of any preceding claim wherein the sample rate of the  
20 samples generated using the motion sensor is higher than the frame rate of  
the video signal.

9. The method of any preceding claim wherein the camera and the motion  
sensor are situated within a mobile device.

10. The method of any preceding claim wherein said using the determined  
25 displacement to compensate for motion in the video signal between the first  
and second frames caused by the motion of the camera comprises:

determining a pixel displacement representing motion in the video  
signal between the first and second frames caused by the determined  
displacement of the camera;

30 filtering the pixel displacement; and

shifting the image of at least one of the first and second frames in  
accordance with the filtered pixel displacement to thereby stabilise the video  
signal.

11. The method of claim 10 wherein said filtering the pixel displacement comprises:

determining an accumulated pixel displacement based on said determined pixel displacement for the second frame; and

5 determining a filtered accumulated pixel displacement for the second frame based on a weighted sum of the determined accumulated pixel displacement for the second frame and a filtered accumulated pixel displacement for the first frame.

10 12. The method of any preceding claim further comprising:

adding a time offset to at least one of (i) the captured plurality of frames, and (ii) the generated plurality of samples, such that the timing of the captured plurality of frames matches the timing of the generated plurality of samples.

15

13. A device for stabilising a video signal, the device comprising:

a camera configured to capture a plurality of frames of the video signal;

a motion sensor, associated with the camera, configured to generate a plurality of samples representing motion of the camera;

20 a displacement determining block configured to use the samples to determine a displacement of the camera between a first time and a second time, wherein the first time corresponds to an exposure time midpoint of a first frame of the video signal and the second time corresponds to an exposure time midpoint of a second frame of the video signal; and

25 a motion compensation block configured to use the determined displacement to compensate for motion in the video signal between the first and second frames caused by the motion of the camera, to thereby stabilise the video signal.

30 14. The device of claim 13 wherein the motion sensor is a gyroscopic motion sensor.

15. The device of claim 13 or 14 wherein the device is a mobile device.

16. A method of stabilising a video signal, the method comprising:  
capturing a plurality of frames of the video signal using a camera;  
determining a pixel displacement representing motion in the video  
5 signal between first and second frames of the video signal caused by motion  
of the camera; and  
shifting an image of at least one of the first and second frames in  
accordance with the pixel displacement to thereby stabilise the video signal,  
wherein the image comprises a stride value indicating memory space of the  
10 image, a plurality of pixel values, a pointer indicating the position of a first  
pixel of the image, and a width value indicating the width of the image,  
wherein said shifting of the image comprises adjusting the pointer and the  
width value without adjusting the stride value.
17. The method of claim 16 wherein the image comprises a plurality of  
15 image planes which are each represented by a respective plurality of pixel  
values, a respective pointer indicating the position of a first pixel of the image  
plane, and a respective width value indicating the width of the image plane,  
wherein the image planes are shifted together by adjusting their respective  
20 pointers and width values without adjusting the stride value.
18. The method of claim 16 or 17 further comprising rounding the  
determined pixel displacement to an integer-pixel displacement.
19. The method of any of claims 16 to 18 further comprising:  
25 using a motion sensor associated with the camera to generate a  
plurality of samples representing motion of the camera; and  
using the samples to determine a displacement of the camera between  
a first frame of the video signal and a second frame of the video signal,  
30 wherein the determined displacement is used to determine said pixel  
displacement.
20. The method of any of claims 16 to 19 further comprising filtering the  
pixel displacement.

21. The method of any of claims 16 to 20 wherein the motion of the camera is rotational motion.

5 22. The method of any of claims 16 to 21 wherein the method is performed in conjunction with the method of any of claims 1 to 12.

23. A device for stabilising a video signal, the device comprising:

a camera configured to capture a plurality of frames of the video signal;

10 a pixel displacement determining block configured to determine a pixel displacement representing motion in the video signal between first and second frames of the video signal caused by motion of the camera; and

an image shifting block configured to shift an image of at least one of the first and second frames in accordance with the pixel displacement to thereby stabilise the video signal, wherein the image comprises a stride value indicating memory space of the image, a plurality of pixel values, a pointer indicating the position of a first pixel of the image, and a width value indicating the width of the image, wherein the image shifting block is configured to shift the image by adjusting the pointer and the width value without adjusting the stride value.

15  
20

24. A method of stabilising a video signal, the method comprising:

capturing a plurality of frames of the video signal using a camera;

25 determining a portion of motion of the camera occurring whilst a shutter of the camera is closed and not whilst the shutter of the camera is open;

determining a pixel displacement representing motion in the video signal between first and second frames of the video signal caused by, and corresponding to, the portion of motion of the camera; and

30 shifting an image of at least one of the first and second frames in accordance with the pixel displacement to thereby stabilise the video signal.

25. The method of claim 24 further comprising:

using a motion sensor associated with the camera to generate a plurality of samples representing the portion of motion of the camera; and

using the samples to determine a displacement of the camera between the first frame of the video signal and the second frame of the video signal,  
5 wherein the determined displacement is used to determine said pixel displacement.

26. The method of claim 25 wherein the motion of the camera is rotational motion, the motion sensor is a rotational motion sensor and the displacement  
10 of the camera is an angular displacement of the camera.

27. The method of any of claims 24 to 26 further comprising filtering the determined pixel displacement.

15 28. The method of any of claims 24 to 27 wherein the method is performed in conjunction with the method of any of claims 1 to 12 or claims 16 to 22.

29. A device for stabilising a video signal, the device comprising:

a camera configured to capture a plurality of frames of the video signal;  
20 a motion determining block configured to determine a portion of motion of the camera occurring whilst a shutter of the camera is closed and not whilst the shutter of the camera is open;

a pixel displacement determining block configured to determine a pixel displacement representing motion in the video signal between first and  
25 second frames of the video signal caused by, and corresponding to, the portion of motion of the camera; and

an image shifting block configured to shift an image of at least one of the first and second frames in accordance with the pixel displacement to thereby stabilise the video signal.  
30

30. A computer program product for stabilising a video signal, the computer program product being embodied on a non-transient computer-readable

medium and configured so as when executed on a processor of a device to perform the operations of any of claims 1 to 12, 16 to 22 or 24 to 28.




Fig 1

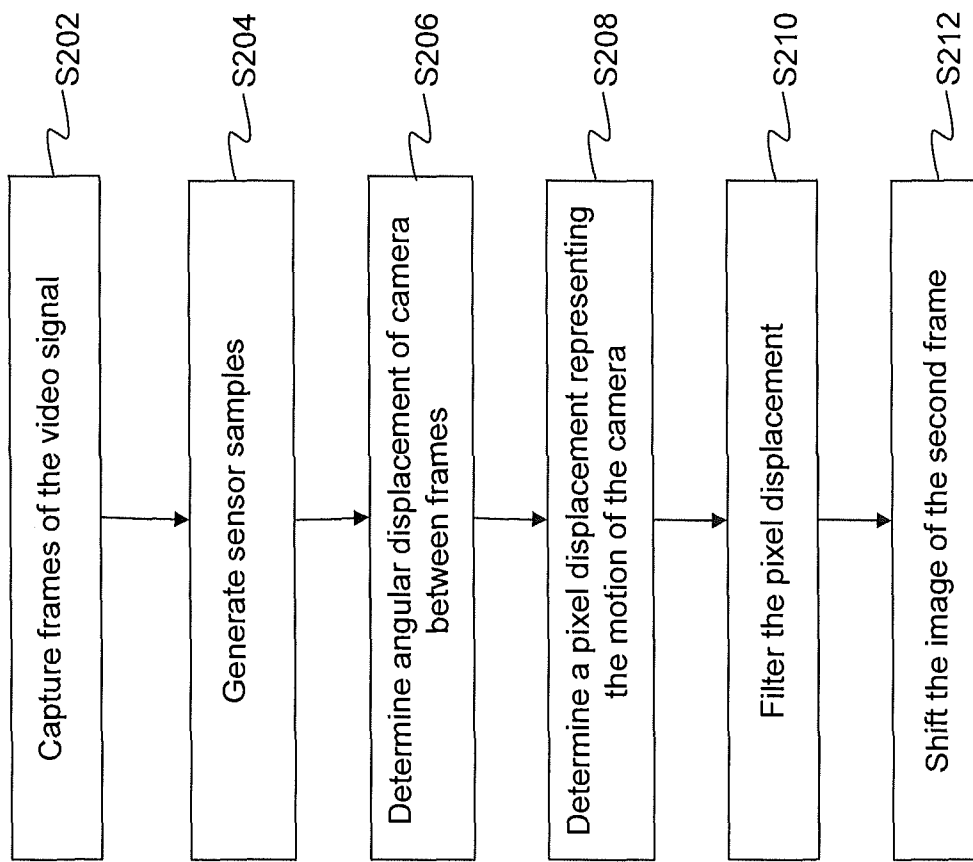


Fig 2




Fig 3




Fig 4




Fig 5