



(51) International Patent Classification:
G06T 1/20 (2006.01)

(21) International Application Number:
PCT/US2014/051799

(22) International Filing Date:
20 August 2014 (20.08.2014)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
14/039,729 27 September 2013 (27.09.2013) US

(71) Applicant: **APPLE INC.** [US/US]; 1 Infinite Loop, Cupertino, California 95014 (US).

(72) Inventors: **ORR, James E.**; 1 Infinite Loop, M/S 74-3PA, Cupertino, California 95014 (US). **MILLET, Timothy John**; 1 Infinite Loop, M/S 74-3PA, Cupertino, California 95014 (US). **CHENG, Joseph J.**; 1 Infinite Loop, M/S 74-3PA, Cupertino, California 95014 (US). **BHARGAVA, Nitin**; 1 Infinite Loop, M/S 85-2SEG, Cupertino, California 95014 (US).

nia 95014 (US). **COTE, Guy**; 1 Infinite Loop, M/S 74-3PA, Cupertino, California 95014 (US).

(74) Agent: **KOWERT, Robert C.**; Meyertons, Hood, Kivlin, Kowert & Goetzel, P.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,

[Continued on next page]

(54) Title: PARALLEL HARDWARE AND SOFTWARE BLOCK PROCESSING PIPELINES




FIG. 3

(57) Abstract: A block processing pipeline that includes a software pipeline and a hardware pipeline that run in parallel. The software pipeline runs at least one block ahead of the hardware pipeline. The stages of the pipeline may each include a hardware pipeline component that performs one or more operations on a current block at the stage. At least one stage of the pipeline may also include a software pipeline component that determines a configuration for the hardware component at the stage of the pipeline for processing a next block while the hardware component is processing the current block. The software pipeline component may determine the configuration according to information related to the next block obtained from an upstream stage of the pipeline. The software pipeline component may also obtain and use information related to a block that was previously processed at the stage.



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).

Published:

- *with international search report (Art. 21(3))*
- *before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments (Rule 48.2(h))*

PARALLEL HARDWARE AND SOFTWARE BLOCK PROCESSING PIPELINES**BACKGROUND****Technical Field**

5 **[0001]** This disclosure relates generally to video or image processing, and more specifically to methods and apparatus for processing digital video frames in block processing pipelines.

Description of the Related Art

10 **[0002]** Various devices including but not limited to personal computer systems, desktop computer systems, laptop and notebook computers, tablet or pad devices, digital cameras, digital video recorders, and mobile phones or smart phones may include software and/or hardware that may implement a video processing method. For example, a device may include an apparatus (e.g., an integrated circuit (IC), such as a system-on-a-chip (SOC), or a subsystem of an IC), that may receive and process digital video input from one or more sources and output the processed
15 video frames according to one or more video processing methods. As another example, a software program may be implemented on a device that may receive and process digital video input from one or more sources and output the processed video frames according to one or more video processing methods. As an example, a video encoder 10 as shown in Fig. 1 represents an apparatus, or alternatively a software program, in which digital video input (input frames 90) is
20 encoded or converted into another format (output frames 92), for example a compressed video format such as H.264 / Advanced Video Coding (AVC) format (also referred to as MPEG 4 Part 10), according to a video encoding method. An apparatus or software program such as a video encoder 10 may include multiple functional components or units, as well as external interfaces to, for example, video input sources and external memory.

25 **[0003]** In some video processing methods, to perform the processing, each input video frame 90 is divided into rows and columns of blocks of pixels (e.g., 16x16 pixel blocks), for example as illustrated in Fig. 2 which shows an example 192x192 pixel frame divided into 144 16x16 pixel blocks. Each block of an input video frame 90 is processed separately, and when done the processed blocks are combined to form the output video frame 92. This may be referred to as a
30 block processing method. Conventionally, the blocks are processed by the block processing method in scan order as shown in Fig. 2, beginning at the first block of the first row of the frame (shown as block 0), sequentially processing the blocks across the row, and continuing at the first block of the next row when a row is complete.

35 **[0004]** A block processing method may include multiple processing steps or operations that are applied sequentially to each block in a video frame. To implement such a block processing

method, an apparatus or software program such as a video encoder 10 may include or implement a block processing pipeline 40. A block processing pipeline 40 may include two or more stages, with each stage implementing one or more of the steps or operations of the block processing method. Fig. 1 shows an example video encoder 10 that implements an example block processing pipeline 40 that includes at least stages 42A through 42C. A block is input to a stage 42A of the pipeline 40, processed according to the operation(s) implemented by the stage 42A, and results are output to the next stage 42B (or as final output by the last stage 42). The next stage 42B processes the block, while a next block is input to the previous stage 42A for processing. Thus, blocks move down the pipeline from stage to stage, with each stage processing one block at a time and multiple stages concurrently processing different blocks. Conventionally, the blocks are input to and processed by the block processing pipeline 40 in scan order as shown in Fig. 2. For example, in Fig. 1, the first block of the first row of the frame shown in Fig. 2 (block 0) is at stage 42C, the second block (block 1) is at stage 42B, and the third block (block 2) is at stage 42A. The next block to be input to the block processing pipeline 40 will be the fourth block in the first row.

H.264 / Advanced Video Coding (AVC)

[0005] H.264/AVC (formally referred to as ITU-T Recommendation H.264, and also referred to as MPEG-4 Part 10) is a block-oriented motion-compensation-based codec standard developed by the ITU-T (International Telecommunications Union – Telecommunication Standardization Sector) Video Coding Experts Group (VCEG) together with the ISO/IEC JTC1 Moving Picture Experts Group (MPEG). The H.264/AVC standard is published by ITU-T in a document titled “ITU-T Recommendation H.264: Advanced video coding for generic audiovisual services”. This document may also be referred to as the H.264 Recommendation.

SUMMARY OF EMBODIMENTS

[0006] Embodiments of block processing methods and apparatus are described in which a block processing pipeline includes a software pipeline and a hardware pipeline that run in parallel. However, the software pipeline runs one or more blocks ahead of the hardware pipeline. The stages of the pipeline may each include a hardware pipeline component that performs one or more operations on a current block at the stage. At least one stage of the pipeline may also include a software pipeline component that determines a configuration for the hardware component at the stage of the pipeline for processing an upcoming block (e.g., the next block) while the hardware component is processing the current block. In at least some embodiments, the software pipeline component at a stage may determine the configuration for processing the

next block at the stage according to information related to the next block obtained from an upstream stage of the pipeline. In at least some embodiments, the software pipeline component may also obtain and use information related to a block that was previously processed at the stage in determining the configuration for processing the next block. In at least some embodiments, the software pipeline component at a stage may obtain information from processing a block or blocks at the stage and pass the information back to one or more upstream stages of the pipeline to configure those upstream stages for processing current or upcoming blocks. In addition, in at least some embodiments, a software pipeline component at or near the end of the pipeline may pass block processing results or statistics back to another software pipeline component at or near the beginning of the pipeline for use in processing upcoming blocks in the pipeline.

[0007] In at least some embodiments, a software pipeline component at a stage receives block information for a next block from an upstream stage. The software pipeline component may also receive information from a block of the frame that was previously processed at the stage. The software pipeline component may determine a configuration for the block according to the received information for the block, write the configuration for the block to a configuration memory, and set a go bit or otherwise signal to the hardware pipeline component at the stage that the configuration for the next block is ready in the configuration memory. The software pipeline component may then push the block information for the block to a downstream stage.

[0008] The hardware pipeline component at the stage receives blocks from a previous stage. If the hardware pipeline component is ready to process the next block, the next block is ready in the memory, and the software pipeline component has signaled to the hardware pipeline component that a configuration for the next block is ready in the configuration memory, then the hardware pipeline component may begin to process the next block. The hardware pipeline component sets the configuration for processing the next block according to the configuration in the configuration memory and signals to the software pipeline component that the configuration memory is available. The hardware pipeline component then processes the block according to the configuration for the block and writes the processed block to the next stage.

[0009] The above operations of the software pipeline component and the hardware pipeline component at a stage in the pipeline may be repeated at the stage until all the blocks in an input frame have been processed at the stage.

BRIEF DESCRIPTION OF THE DRAWINGS

[0010] Fig. 1 illustrates an example video encoder including a conventional block processing pipeline that processes blocks from input frames in scan order.

[0011] Fig. 2 illustrates conventional scan order processing of blocks from a video frame.

[0012] Fig. 3 is a high-level block diagram of an example block processing pipeline that implements a software pipeline and a hardware pipeline, according to at least some embodiments.

5 [0013] Figs. 4A through 4C illustrate processing blocks at a stage in an example block processing pipeline that implements a software pipeline and a hardware pipeline, according to at least some embodiments.

[0014] Fig. 5 illustrates an example block processing pipeline that implements a software pipeline and a hardware pipeline in which at least one stage is skipped by the software pipeline, according to at least some embodiments.

10 [0015] Fig. 6 illustrates an example block processing pipeline that implements a software pipeline and a hardware pipeline in which at least one stage includes multiple pipeline units, according to at least some embodiments.

[0016] Fig. 7 illustrates components of an example pipeline unit that may be used at a stage of a block processing pipeline that implements a software pipeline and a hardware pipeline, according to at least some embodiments.

15 [0017] Figs. 8A and 8B are flowcharts of methods of operation of a software pipeline and a hardware pipeline that operate in parallel in a block processing pipeline, according to at least some embodiments.

[0018] Fig. 9 shows neighbor blocks of a current block in a frame, and further illustrates a knight's order processing method for the blocks, according to at least some embodiments.

20 [0019] Figs. 10A and 10B graphically illustrate the knight's order processing method including the algorithm for determining a next block, according to at least some embodiments.

[0020] Figs. 11A and 11B are high-level flowcharts of a knight's order processing method for a block processing pipeline, according to at least some embodiments.

25 [0021] Fig. 12 shows a portion of a quadrow as processed in a pipeline according to the knight's order processing method that may be cached in the current quadrow buffer, according to at least some embodiments

[0022] Fig. 13 graphically illustrates blocks in a current quadrow being processed according to the knight's order processing method, as well as neighbor blocks in the last row of the previous quadrow that may be cached in a previous quadrow buffer, according to at least some

30 [0023] Fig. 14 is a flowchart of a method for processing blocks in a block processing pipeline in which neighbor data is cached in local buffers at the stages of the pipeline, according to at least some embodiments.

[0024] Figs. 15A and 15B are block diagrams of example pipeline processing units that may be used at the stages of a block processing pipeline that implements one or more of the block processing methods and apparatus as described herein, according to at least some embodiments.

5 [0025] Fig. 15C shows that a single processor may be associated with a group of two or more pipeline units.

[0026] Fig. 16 is a high-level block diagram of general operations in an example block processing method that may be implemented by a block processing pipeline that implements one or more of the block processing methods and apparatus described herein, according to at least some embodiments.

10 [0027] Fig. 17 is a block diagram of an example video encoder apparatus, according to at least some embodiments.

[0028] Fig. 18 is a block diagram of one embodiment of a system on a chip (SOC).

[0029] Fig. 19 is a block diagram of one embodiment of a system.

[0030] While the invention is susceptible to various modifications and alternative forms, 15 specific embodiments thereof are shown by way of example in the drawings and will herein be described in detail. It should be understood, however, that the drawings and detailed description thereto are not intended to limit the invention 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 present invention as defined by the appended claims. As used throughout this application, the word “may” is used in a permissive sense (i.e., meaning having the potential 20 to), rather than the mandatory sense (i.e., meaning must). Similarly, the words “include,” “including,” and “includes” mean including, but not limited to.

[0031] Various units, circuits, or other components may be described as “configured to” perform a task or tasks. In such contexts, “configured to” is a broad recitation of structure 25 generally meaning “having circuitry that” performs the task or tasks during operation. As such, the unit/circuit/component can be configured to perform the task even when the unit/circuit/component is not currently on. In general, the circuitry that forms the structure corresponding to “configured to” may include hardware circuits. Similarly, various units/circuits/components may be described as performing a task or tasks, for convenience in the 30 description. Such descriptions should be interpreted as including the phrase “configured to.” Reciting a unit/circuit/component that is configured to perform one or more tasks is expressly intended not to invoke 35 U.S.C. § 112, paragraph six, interpretation for that unit/circuit/component.

DETAILED DESCRIPTION

[0032] In the following description, numerous specific details are set forth to provide a thorough understanding of the present invention. However, one having ordinary skill in the art should recognize that the invention might be practiced without these specific details. In some instances, well-known circuits, structures, and techniques have not been shown in detail to avoid obscuring the present invention.

[0033] Various embodiments of methods and apparatus for processing digital video frames in block processing pipelines are described. Embodiments of block processing methods and apparatus are generally described herein in the context of video processing in which input video frames are subdivided into and processed according to blocks of elements (e.g., 16x16, 32x32, or 64x64 pixel blocks). Embodiments of an example H.264 video encoder that includes a block processing pipeline and that may implement one or more of the block processing methods and apparatus are described herein. The H.264 video encoder converts input video frames from an input format into H.264 / Advanced Video Coding (AVC) format as described in the H.264/AVC standard (the H.264 Recommendation). Fig. 16 illustrates an example block processing pipeline of an example H.264 video encoder, and Fig. 17 illustrates an example H.264 video encoder that includes a block processing pipeline. However, embodiments of the block processing methods and apparatus may be used in encoders for other video encoding formats, for example in block processing pipelines of HEVC (High Efficiency Video Encoding) video encoders that convert input video frames from an input format into HEVC format as described in the HEVC standard. Other video encoders that may use embodiments of the block processing methods and apparatus may include, but are not limited to, H.263, MPEG-2, MPEG-4, and JPEG-2000 video encoders. However, it is to be noted that embodiments of the block processing methods and apparatus may be used in any block processing pipeline, including but not limited to block processing pipelines implemented in various other video encoders and/or decoders (which may be referred to as codecs) in which digital video frames input in one format are encoded or converted into another format. Further note that the block processing methods and apparatus may be used in software and/or hardware implementations of video encoders. In addition to video encoders/decoders, the block processing methods and apparatus described herein may be used in various other applications in which blocks from a video frame or still digital image are processed, for example in pipelines that process still digital images in various image processing applications. Thus, it is to be understood that the term frame or video frame as used herein may also be taken to refer to any digital image.

[0034] Embodiments of the block processing methods and apparatus as described herein may be implemented in two or more parallel block processing pipelines. For example, 2, 4, 8, or

more pipelines may be configured to run in parallel, with each pipeline processing a quadrow from an input video frame, for example with blocks input according to knight's order.

[0035] Embodiments of the block processing methods and apparatus are generally described herein in the context of video processing in which input frames are subdivided into and
5 processed according to blocks of picture elements (referred to as pixels, or pels), specifically 16x16 pixel blocks referred to as macroblocks that are used, for example, in H.264 encoding. However, embodiments may be applied in pipelines in which blocks of other sizes and geometries, or of other elements, are processed. For example, HEVC encoding uses blocks referred to as Coding Tree Units (CTUs) that may vary within the range of 16x16 pixel to 64x64
10 pixel. In some implementations such as H.264 encoders, the blocks input to the pipeline may be referred to as macroblocks, each macroblock including two or more blocks or partitions that may be processed separately at stages of the pipeline. For example, for input video frames encoded in YUV (e.g., YUV420 format) or YCbCr (e.g., YCbCr 4:2:0, 4:2:2 or 4:4:4 formats) color space, a macroblock may be composed of separate blocks of chroma and luma elements that may be
15 processed separately at stages in a pipeline. In addition to applications that process frames in a pipeline according to blocks of elements (e.g., blocks of pixels), the block processing methods and apparatus may be applied in applications in which digital images (e.g., video frames or still images) are processed by single elements (e.g., single pixels).

20 Parallel hardware and software block processing pipelines

[0036] Embodiments of block processing methods and apparatus are described in which a block processing pipeline includes a software pipeline and a hardware pipeline that run in parallel. However, the software pipeline runs one or more blocks ahead of the hardware pipeline.

The stages of the pipeline may each include a hardware pipeline component (e.g., a circuit) that
25 performs one or more operations on a current block at the stage. At least one stage of the pipeline may also include a software pipeline component that determines a configuration for the hardware component at the stage of the pipeline for processing an upcoming block (e.g., the next block) while the hardware component is processing the current block. The software pipeline component may include at least a processor. In at least some embodiments, the software pipeline
30 component at a stage may determine the configuration for processing the next block at the stage according to information related to the next block obtained from an upstream stage of the pipeline. In at least some embodiments, the software pipeline component may also obtain and use information related to a block that was previously processed at the stage in determining the configuration for processing the next block. In at least some embodiments, the software pipeline
35 may also "look ahead" (upstream) one or more blocks to obtain information from upcoming

blocks that may be used in determining the configurations for processing the next blocks at the stages. The software pipeline components may generate statistics on one or more blocks that are used in determining the configurations. In at least some embodiments, the software pipeline component at a stage may obtain information from processing a block or blocks at the stage and pass the information back to one or more upstream stages of the pipeline to configure those upstream stages for processing current or upcoming blocks. In addition, in at least some embodiments, a software pipeline component at or near the end of the pipeline may pass block processing results or statistics back to another software pipeline component at or near the beginning of the pipeline for use in processing upcoming blocks in the pipeline.

[0037] The block information obtained by a software pipeline component at a stage and used to determine a configuration for processing a next block at the stage may, for example, include various statistics related to the block and/or to one or more other blocks. The following provides some examples of block statistics that may be used in some embodiments, and is not intended to be limiting:

- Sum of pixels (s).
- Sum of pixels squared (s^2).
- Block variance (may be estimated from s and s^2 , e.g. $var = s^2 - (s)^2$).
- Horizontal and vertical gradients (G_x and G_y).
- Gradient histograms for G_x and G_y .

[0038] The operations performed by the hardware pipeline components at the various stages may vary, and thus the configuration for the hardware pipeline components at the stages may vary. Thus, the software pipeline components at the stages may determine and set particular configuration parameters according to the respective hardware pipeline components at the stages. However, a general example of configuration parameters that may be determined and set at a stage by the software pipeline component based on an analysis of the information is given below, and is not intended to be limiting.

[0039] One or more stages of a pipeline may perform operations to determine a best mode for processing pixels in a given block. At a particular stage, the hardware pipeline component may receive information from one or more upstream stages (and possibly feedback from one or more downstream stages) and use this information to select a particular one of multiple modes. The software pipeline component at the stage may receive, generate, and analyze statistics related to the block (e.g., block variance) and set one or more configuration parameters according to the analysis to, for example, cause the hardware pipeline component to try multiple modes if

the block variance is high, or to bias the hardware component towards a particular mode or modes if the block variance is low.

[0040] In at least some embodiments, a block processing pipeline that implements parallel software and hardware pipelines may input blocks to and process blocks in the pipelines according to knight's order, as described in the section titled Knight's order processing. However, other block input and processing orders may be used in some embodiments. In at least some embodiments, at least one stage of a block processing pipeline that implements parallel software and hardware pipelines may implement one or more local buffers for caching data for neighbor blocks at the stage, as described in the section titled Caching neighbor data.

[0041] Fig. 3 is a high-level block diagram of an example block processing pipeline 300 that implements a software pipeline 302 and a hardware pipeline 304, according to at least some embodiments. The software pipeline 302 and the hardware pipeline 304 process blocks from a frame in parallel, with the software pipeline 302 at least one block ahead of the hardware pipeline 304. The pipeline 300 may include multiple stages 320, each stage configured to perform one or more operations on a block of pixels from a frame (e.g., a video frame). At least some of the stages (stages 320A and 320B in Fig. 3) may each include at least one pipeline unit 330 that includes a software pipeline component 322 and a hardware pipeline component 326. The hardware pipeline component 326 of each pipeline unit 330 may perform one or more particular operations of a block processing method on a block currently at the stage 320 in the hardware pipeline 304. While the hardware pipeline component 326 of a given pipeline unit 330 is working on the current block at the stage 320, the software pipeline component 322 of the pipeline unit 330 at the stage 320 may preconfigure the hardware pipeline component 326 for processing an upcoming block, which may be the next block. Thus, the software pipeline 302 operates one or more blocks ahead of the hardware pipeline 304.

[0042] For example, as shown in Fig. 3, at stage 320B hardware pipeline component 326B is currently processing block i while software pipeline component 326B is configuring the hardware pipeline component 326B to process block $i+1$, and at stage 320A hardware pipeline component 326A is currently processing block $i+1$ while software pipeline component 326A is configuring the hardware pipeline component 326A to process block $i+2$.

[0043] The software pipeline component 322 of a pipeline unit 330 at a stage 320 may determine a configuration for processing an upcoming block (e.g., a next block) at the hardware pipeline component 326 of the respective pipeline unit 330 according to information for the block. The information for the block may include at least block information received from an upstream stage. In at least some embodiments, the information may also include feedback information from one or more blocks previously processed at the stage 320. The software

pipeline component 322 may preconfigure the hardware pipeline component 326 of the pipeline unit 330 at the stage 320 for processing the block according to the determined configuration, for example by setting one or more configuration values in a set of registers or other memory coupled to the hardware pipeline component 326. Once the configuration for processing the block at the hardware pipeline component 326 of the pipeline unit 330 is ready, the software pipeline component 322 may signal the hardware pipeline component 326 of the pipeline unit 330. Assuming that the hardware pipeline component 326 has completed the processing of a previous block and that the next block is available to the hardware pipeline component 326 (e.g., ready to be read from its input buffer), the hardware pipeline component 326 of the pipeline unit 330 may then begin processing the next block according to the configuration for the block that was determined and preconfigured by the software pipeline component 322 of the pipeline unit 330.

[0044] In at least some embodiments, an initial stage 310 of the pipeline may input block information to the software pipeline 302 and blocks to the hardware pipeline 304. The initial stage 310 may obtain block input, for example from an external memory via direct memory access (DMA), and buffer the blocks in a block buffer component 312. Block buffer component 312 may have the capacity to hold one, two, or more blocks. For example, in some embodiments, block buffer component 312 may be able to buffer 16 blocks. In at least some embodiments, block buffer component 312 may buffer one, two or more blocks for input to the hardware pipeline 304 before initial stage 310 begins input of blocks to the hardware pipeline 304. In at least some embodiments, once the initial stage 310 begins input of blocks to the hardware pipeline 304, the initial stage 310 may write a next block from block buffer component 312 to a buffer memory of the hardware pipeline component 326A of pipeline unit 330A at stage 320A when the pipeline unit 330A is ready to receive the next block. The initial stage 310 may continue to obtain block input for a frame, buffer the blocks to block buffer component 312, and input blocks to the hardware pipeline 304 until all the blocks in the frame are processed.

[0045] A block analysis component 314 at initial stage 310 may perform one or more analysis functions on one or more blocks that are currently buffered in block buffer component 312 including a next block to be input to the hardware pipeline 304 to generate block information for the next block. The block information may, for example, include one or more block statistics. Some non-limiting examples of block statistics that may be generated were previously provided. Once the block information is generated for the next block, the initial stage 310 may send the block information to the software pipeline component 322A of the pipeline unit 330A at stage 320A of the pipeline 300. The block analysis component 314 may continue to generate block

information and input the block information to the software pipeline 304 until all the blocks in the frame are processed.

[0046] In at least some embodiments, the software pipeline component 322 of each pipeline unit 330 may include a memory for buffering block information for one, two, or more upcoming blocks. In at least some embodiments, the hardware pipeline component 326 of each pipeline unit 330 may include a memory for storing one or more blocks to be processed at the stage 320. In at least some embodiments, the memory may be a double buffer so that a previous stage can write a next block to the memory while the hardware pipeline component 326 is reading a current block from the memory.

[0047] In at least some embodiments, the software pipeline component 322 of a pipeline unit 330 may push block information for each block to the software pipeline component 322 of a pipeline unit 330 at a downstream stage 320 so that the software pipeline component 322 at the downstream stage 320 can configure the respective hardware pipeline component 326 at the stage. In at least some embodiments, the software pipeline component 322 of a pipeline unit 330 at a stage 320 does not push block information for a block to a downstream stage 320 until after completing the preconfiguration for processing the block at the stage 320. In at least some embodiments, the block information for a block may be updated according to information that is available at a stage 320 before pushing the block information to the downstream stage 320.

[0048] Once a hardware pipeline component 326 at a stage 320 has completed processing of a block, the processed block may be sent to a hardware pipeline component 326 at the next stage 320 for processing. The hardware pipeline component 326 at the next stage 320 may hold the block in its memory until the hardware pipeline component 326 has completed processing of a current block and has received a signal from the software pipeline component 322 of the pipeline unit 330 at the stage 320 that the configuration for processing the block is ready. Note that a processed block may instead be written to a memory external to the pipeline 300 by a last stage 320 of the pipeline 300.

[0049] Figs. 4A through 4C illustrate processing blocks at a pipeline unit of a stage in an example block processing pipeline that implements a software pipeline and a hardware pipeline, according to at least some embodiments. Figs. 4A through 4C show a pipeline unit 330 that may be used at a stage in a block processing pipeline that includes a software pipeline component 322 and a hardware pipeline component 326. The hardware pipeline component 326 of the pipeline unit 330 may perform one or more particular operations of a block processing method on a block currently at the stage in the hardware pipeline 304. While the hardware pipeline component 326 is working on the current block, the software pipeline component 322 of pipeline unit 330 may preconfigure the hardware pipeline component 326 for processing an upcoming block, for

example the next block. Thus, the software pipeline component 322 of a pipeline unit 330 operates at least one block ahead of the hardware pipeline component 326 of the pipeline unit 330.

[0050] The pipeline unit 330 may also include a configuration memory (shown as config memory 324A and 324B in Figs. 4A through 4C). The configuration memory may, for example, be a set of hardware registers. As shown in Figs. 4A through 4C, in at least some embodiments, the configuration memory may be partitioned into two memories (config memory 324A and 324B) so that the software pipeline component 322 of pipeline unit 330 can write to one memory while the hardware pipeline component 326 is reading from the other memory. The configuration memory may, for example, be a set of registers that are partitioned into a subset of active registers to which the software pipeline component 322 writes the configuration for a next block and a subset of shadow registers from which the hardware pipeline component 326 reads the configuration for a current block. In at least some embodiments, the software pipeline component 322 may write to either of the config memories 324A and 324B, and the hardware pipeline component 326 may read from either of the config memories 324A and 324B; the two components may both toggle between the memories 324, with the software pipeline component 322 writing to one while the hardware pipeline component 326 is reading from the other. Alternatively, in some embodiments, the software pipeline component 322 may write to only one of the config memories 324 (e.g., config memory 324A), and the hardware pipeline component 326 may read from only the other config memory 324 (e.g., config memory 324B); when the hardware pipeline component 326 is ready for a new configuration and the configuration is ready, the configuration may be copied from the config memory 324A to the config memory 324B. Note that embodiments may also be implemented in which only a single configuration memory is used, or in which more than two configuration memories are used.

[0051] Fig. 4A show a pipeline unit 330 of a stage at an initial state. Software pipeline component 322 receives, from an upstream stage, block information for a first block (block *i*) from a frame to be processed at the stage. Hardware pipeline component 326 is not currently processing a block. Software pipeline component 322 determines a configuration for processing block *i* according to the received block information and writes the configuration to config memory 324A. Software pipeline component 322 signals hardware pipeline component 326 of pipeline unit 330 that the configuration for block *i* is ready, for example by setting a go bit or flag.

[0052] Fig. 4B show the pipeline unit 330 at the next cycle. Software pipeline component 322 pushes block information for block *i* to a downstream stage. Hardware pipeline component 326 receives block *i* and processes block *i* according to the configuration in config memory

324A. Software pipeline component 322 receives block information for a next block (block $i+1$) to be processed at the stage. Software pipeline component 322 determines a configuration for processing block $i+1$ according to the received block information and writes the configuration to config memory 324B. Software pipeline component 322 signals hardware pipeline component 326 that the configuration for block $i+1$ is ready, for example by setting a go bit or flag.

[0053] Fig. 4C shows the pipeline unit 330 at the next cycle. Software pipeline component 322 pushes block information for block $i+1$ to a downstream stage. Hardware pipeline component 326 receives block $i+1$ and processes block $i+1$ according to the configuration in config memory 324B. Software pipeline component 322 receives block information for a next block (block $i+2$) to be processed at the stage. Software pipeline component 322 determines a configuration for processing block $i+2$ according to the received block information and writes the configuration to config memory 324A. Software pipeline component 322 signals hardware pipeline component 326 that the configuration for block $i+2$ is ready, for example by setting a go bit or flag.

[0054] Fig. 4C also shows that information from a previously processed block at a stage may be obtained by the software pipeline component 322 at the stage and used in determining a configuration for a next block to be processed by the hardware pipeline component 326 at the stage. Hardware pipeline component 326 finished processing block i at a previous cycle, as shown in Fig. 4B, and is now processing block $i+1$ at Fig. 4C. Thus, information from the processing of block i at the stage is available, and may be fed back to the software pipeline component 322 of the pipeline unit 330 at the stage. This information from the processing of block i at the stage may be used in combination with the block information for block $i+2$ received from an upstream stage to determine the configuration for block $i+2$. Thus, feedback of information from the processing of blocks at a stage may be for a block that is two ahead of the block for which a configuration is being generated.

[0055] Alternatively, in some implementations, the software pipeline component 322 may wait for completion of the processing of a current block by the hardware pipeline component 326 at the stage, and use this information to determine a configuration for the next block. In this case, feedback of information from the processing of blocks at a stage may be for a block that is only one ahead of the block for which a configuration is being generated

[0056] Fig. 5 illustrates an example block processing pipeline 300 that implements a software pipeline and a hardware pipeline in which at least one stage is skipped by the software pipeline, according to at least some embodiments. In some pipeline implementations, one or more pipeline units 330 of the pipeline 300 may include a hardware pipeline component 326 that does not require dynamic configuration. Fig. 5 shows three stages 320A, 320B, and 320C. Stage

320A includes pipeline unit 330A that includes both a software pipeline component 322A and a hardware pipeline component 326A, and stage 320C includes a pipeline unit 330C that includes both a software pipeline component 322B and a hardware pipeline component 326C. However, stage 320B includes a pipeline unit 330B that includes a hardware pipeline component 326B that
5 does not require dynamic configuration, as the operation(s) the component 326 performs on a block are the same for all blocks. Thus, pipeline unit 330B does not utilize a software pipeline component 322.

[0057] As shown in Fig. 5, hardware pipeline component 326A at stage 320A is currently processing block $i+2$, while software pipeline component 322A at stage 320A is determining and
10 setting the configuration for processing the next block ($i+3$) at stage 320A. Hardware pipeline component 326B at stage 320B is currently processing block $i+1$. Hardware pipeline component 326C at stage 320C is currently processing block i , while software pipeline component 322B at stage 320C is determining and setting the configuration for processing the next block ($i+1$) at stage 320A. In at least some embodiments, the block information for block $i+2$ may be pushed
15 downstream from software pipeline component 322A to software pipeline component 322B once stage 320A completes the configuration for processing block $i+2$ and buffered at software pipeline component 322B until software pipeline component 322B is ready to configure hardware pipeline component 322C to process block $i+2$. Alternatively, stage 320B may include buffers to which block information is pushed from stage 320A and from which block information
20 is pushed to stage 320C. As another alternative, stage 320A may buffer block information that it is done with until stage 320C is ready for the information.

[0058] Fig. 6 illustrates an example block processing pipeline 300 that implements a software pipeline and a hardware pipeline in which at least one stage includes multiple pipeline units, according to at least some embodiments. As shown in Fig. 6, stage 320A includes a single
25 pipeline unit 330A that includes a software pipeline component 322A and a hardware pipeline component 326A, and stage 320C includes a single pipeline unit 330C that includes a software pipeline component 322C and a hardware pipeline component 326D. However, stage 320B include two pipeline units 320B and 320C. Pipeline unit 330B includes a software pipeline component 322B and a hardware pipeline component 326B. Pipeline unit 330C includes only a
30 hardware pipeline component 326C. In hardware pipeline 304, blocks or portions of blocks from pipeline unit 330A at stage 320A pass through both hardware pipeline component 326B and hardware pipeline component 326C of stage 320B, which output processed blocks or portions of blocks to hardware pipeline component 326D of pipeline unit 330D in stage 320C. In software pipeline 302, block information is passed from software pipeline unit 322A at stage 320A to

software pipeline unit 322B at stage 320B, and from software pipeline unit 322B at stage 320B to software pipeline unit 322C at stage 320C.

[0059] While not shown, in some implementations, a stage may include two or more pipeline units 330 that include both a software pipeline component 322 and a hardware pipeline component 336. In this case, an upstream stage may feed block information to the software pipeline component 322 of each pipeline unit at the stage 320. However, in at least some embodiments, only one of the software pipeline components 322 may push the block information to a software pipeline component 322 of a pipeline unit 330 at a downstream stage 320.

[0060] Fig. 7 illustrates components of an example pipeline unit that may be used at a stage of a block processing pipeline that implements a software pipeline and a hardware pipeline, according to at least some embodiments. As shown in Fig. 7, the hardware pipeline component 404 of a pipeline unit 400 may include at least a memory 432 and a unit core 430. Unit core 5020 may be a component (e.g., a circuit) that is configured to perform a particular operation on or for a block, or a portion of a block, at a particular stage of the block processing pipeline. Memory 432 may, for example, be a double-buffered memory that allows the unit core 430 to read and process data for a block from the memory 432 while data for a next block is being written to the memory 432 from a previous pipeline unit.

[0061] As shown in Fig. 7, a pipeline unit 400, in addition to a hardware pipeline component 404 that includes memory 432 and unit core 430, may also include a software pipeline component 402 that includes at least a processor 410 and a memory 412. Processor 410 may, for example, be a mobile or M-class processor. The processors 410 may, for example, be configured to determine and set configurations for a next block to be processed at the hardware pipeline unit 404 according to block information received at the software pipeline component 402. In at least some embodiments, the processor 410 may also be configurable, for example with low-level firmware microcode, to allow flexibility in algorithms that are implemented by the block processing pipeline for various applications.

[0062] In at least some embodiments, the software pipeline component 402 may be configured to receive block information from a previous (upstream) stage of the pipeline and send block information to a subsequent (downstream) stage of the pipeline. In addition, a software pipeline component 402 at a last stage of the pipeline may be configured to send feedback data to an upstream stage (e.g. the first stage) of the pipeline. In at least some embodiments, the software pipeline component 402 may also receive information for a block that was previously processed by the hardware pipeline component 404 of the pipeline unit 400.

[0063] Software pipeline component 402 may buffer block information received from an upstream stage of the pipeline in memory 412, and push block information from memory 412 to a downstream stage of the pipeline. In at least some embodiments, memory 412 may be a double buffer memory so that an upstream stage can push block information for a next block to the software pipeline component 402 while the processor 410 is accessing block information for a previous block from the memory 412. In some embodiments, memory 412 may be able to buffer more than two sets of block information, for example in cases where the previous stage does not include a software pipeline component as shown by stage 320B in Fig. 5.

[0064] The processors 410 may read block information for a next block from memory 412 and determine a configuration for the next block according to the block information. In at least some embodiments, the processor 410 may also receive information for a block that was previously processed by the hardware pipeline component 404 of the pipeline unit 400 and use that information in determining the configuration for the next block.

[0065] As shown in Fig.7, a pipeline unit 400 may also include an interface 406 between software pipeline component 402 and hardware pipeline component 404. In at least some embodiments, the interface 406 may be a set of registers. Note, however, that the interface 406 may be otherwise implemented. In the pipeline unit 400 as shown in Fig. 7, interface 406 includes at least config memory 420A, config memory 420B, and go 422. In at least some embodiments, the processor 410 may write to either of the config memories 420A and 420B, and the unit core 430 may read from either of the config memories 420A and 420B; the processor 410 and unit core 430 may toggle between the two memories 420, with the processor 410 writing to one while the unit core 430 is reading from the other. Alternatively, in some embodiments, the processor 410 may write to only one of the config memories 420 (e.g., config memory 420A), and the unit core 430 may read from only the other config memory 420 (e.g., config memory 420B); when the unit core 430 is ready for a new configuration and the configuration is ready, the configuration may be copied from config memory 420A to config memory 420B. Note that embodiments may also be implemented in which only a single configuration memory is used, or in which more than two configuration memories are used.

[0066] Go 422 may, for example, be implemented as one or more bits in a register or other memory, or may be otherwise implemented. In at least some embodiments, when processor 410 completes a configuration for a next block and has set the config memory 420 (e.g., config memory 420A) with the configuration, processor 410 may set go 422 to signal to the unit core 430 that the configuration for the next block is ready in the config memory 420 (e.g., config memory 420A). Unit core 430 may begin to process the next block once go 422 is set for the next block if processing of the current block is complete and the next block is ready in memory

432. Otherwise, unit core 430 may wait until completion of processing of the current block and/or the next block is ready in memory 432. Note that initially (for a first block in the pipeline) no block is being processed at the stage when the processor 410 configures the stage for the first block, and thus unit core 430 may begin to process the first block once go 422 is set for the first block and the first block is ready in memory 432. Once unit core 430 is done with the configuration in a config memory 420, the unit core 430 may clear go 422 to signal to processor 410 that the config memory 420 is available to receive the configuration for a next block.

[0067] Figs. 8A and 8B are flowcharts of methods of operation of a software pipeline and a hardware pipeline that operate in parallel in a block processing pipeline to process the blocks from a frame, according to at least some embodiments. Fig. 8A shows operations at a stage for the software pipeline, and Fig. 8B shows operations at the stage for the hardware pipeline. Note that the software pipeline runs at least one block ahead of the hardware pipeline.

[0068] Referring to Fig. 8A, as indicated at 500, a software pipeline component at a stage receives block information. The block information may include block information received from an upstream stage, and/or block information received from a downstream stage. In at least some embodiments, the software pipeline component may also receive information from a block of the frame that was previously processed at the stage. The block information for one, two, or more blocks may be buffered in a local memory of the software pipeline component. As indicated by the arrow that returns to element 500, element 500 may be iteratively performed as long as there are blocks in the frame to be processed.

[0069] As indicated at 502, once the block information for a next block is ready at the stage, the software pipeline component may determine a configuration for the block according to the received information for the block. As indicated at 504, the software pipeline component may write the configuration for the block to a configuration memory of the stage. As indicated at 506, the software pipeline component may set a go bit or otherwise signal to the hardware pipeline component at the stage that the configuration for the next block is ready in the configuration memory. As indicated at 506, the software pipeline component may then push the block information for the block to a downstream stage. At 510, if there are more blocks from the frame to be processed at the stage, the software pipeline method may return to element 502 to begin configuring the hardware pipeline component for a next block. Otherwise, processing of the frame at this stage is done, and the method completes.

[0070] Referring to Fig. 8B, a hardware pipeline component at a stage receives blocks to be processed from a previous stage. The block information may be buffered in a local memory of the hardware pipeline component. In at least some embodiments, the local memory may be a double buffer memory so that the previous stage can write a next block to the stage while the

hardware pipeline component is processing a current block from the memory. As indicated by the arrow that returns to element 550, element 550 may be iteratively performed as long as there are blocks in the frame to be processed.

[0071] At 552, if the hardware pipeline component is not currently processing a block, a next block is ready in the memory, and the software pipeline component has signaled to the hardware pipeline component that a configuration for the next block is ready in the configuration memory (e.g., by setting a go bit or flag), then the hardware pipeline component may begin to process the next block. In at least some embodiments, if any of these three conditions is not met, the hardware pipeline component waits until all three are met. Note, however, that for a first block in the frame to be processed, there will not be a current block being processed at the hardware pipeline component when the first block is received for processing at the hardware pipeline component.

[0072] If all necessary conditions are met, then as indicated 554 the hardware pipeline component sets the configuration for processing the next block according to the configuration in the configuration memory. As indicated at 556, the hardware pipeline component clears the go bit or otherwise signals to the software pipeline component that the configuration memory is available. As indicated at 558, the hardware pipeline component processes the block according to the configuration for the block. As indicated at 560, the hardware pipeline component writes the processed block to the next stage. Alternatively, at a last stage, the processed block may be written to a memory, for example to an external memory via direct memory access (DMA). At 562, if there are more blocks from the frame to be processed at the stage, the hardware pipeline method may return to element 552 to begin processing a next block when all conditions are met. Otherwise, processing of the frame at this stage is done, and the method completes.

[0073] Note that elements 502 through 508 of Fig. 8A are performed by the software pipeline component at a stage for an initial block in the frame before elements 554 through 560 of Fig. 8B are performed by the hardware pipeline component at the stage. After that, elements 502 through 508 of Fig. 8A may be performed by the software pipeline component at the stage to configure the hardware pipeline component for a next block while elements 554 through 560 of Fig. 8B are performed by the hardware pipeline component to process a current block.

Knight's order processing

[0074] Embodiments of block processing methods and apparatus are described in which, rather than processing blocks in a pipeline according to scan order as in conventional methods, the blocks are input to and processed in the pipeline according to an order referred to herein as "knight's order." Knight's order is in reference to a move of a chess knight piece in which the

knight moves one row down and two columns to the left. Note, however, that “knight’s order” as used herein more generally encompasses movements of one row down and p columns to the left, where p may be but is not necessarily 2.

[0075] The knight’s order processing method may provide spacing (one or more stages) between adjacent blocks in the pipeline, which, for example, facilitates feedback of data from a downstream stage of the pipeline processing a first block to an upstream stage of the pipeline processing a second block that depends on the data from the first block. One or more stages of a block processing pipeline may require information from one or more other neighbor blocks when processing a given block. Fig. 9 shows neighbors of a current block (m,n) from which information may be required – left $(m-1,n)$; top $(m,n-1)$; top-left $(m-1,n-1)$; top-right $(m+1,n-1)$; and top-right-right $(m+2,n-1)$. These requirements for information from neighbor block(s) may be referred to as dependencies. For example, referring to Fig. 9, information from the left neighbor of block (m,n) may be required to perform a particular operation on the block. In the knight’s order processing method, rather than inputting block $(m+1, n)$ into the pipeline immediately after block (m,n) , the next block input to the pipeline is block $(m-2,n+1)$. Inputting the blocks into the pipeline in knight’s order rather than scan order provides spacing (e.g., one or more stages) between adjacent blocks on a row in the pipeline.

[0076] In at least some embodiments of the knight’s order processing method, the rows of blocks in the input frame may be divided into sets of four rows, referred to herein as quadrows, with the knight’s order processing method constrained by the quadrow boundaries. Referring to Fig. 9, when using quadrow boundaries with knight’s order processing block $(m-1,n)$ will be four stages downstream when block (m,n) is input to the pipeline, and block (m,n) will be four stages downstream when block $(m+1,n)$ is input to the pipeline. Thus, blocks that are adjacent on a row will be spaced four stages apart in the pipeline. Thus, at stages in which operations are performed on a block that depend on left neighbor information, the information for the left neighbor is more likely to be readily available with less latency than it would be if processing the blocks in scan order. In addition to dependencies on the left neighbor, one or more operations of a block processing method may depend on neighbor blocks from the previous (or above) row such as the top neighbor, top-left neighbor, top-right neighbor, and top-right-right neighbor blocks as shown in Fig. 9. The knight’s order processing method with quadrow constraints provides locality of neighbor information that may be leveraged to provide local caching of this neighbor data at each stage in relatively small buffers.

[0077] In at least some embodiments, a basic algorithm for determining a next block to input to the pipeline according to the knight’s order processing method using quadrow constraints is as follows:

If not on the bottom row of a quadrow:

The next block is two columns left, one row down $(-2,+1)$.

Otherwise, at the bottom row of a quadrow:

The next block is seven columns right, three rows up $(+7,-3)$.

5

[0078] However, the knight's order processing method may also be implemented with other spacing than two blocks left, one block down $(-2,+1)$. For example, instead of two blocks left and one block down, the method may be implemented to go three blocks left and one block down to get the next block. As another example, the method may be implemented to go one block left and one block down $(-1,+1)$ to get the next block. In addition, the knight's order processing method may be implemented with other row constraints than quadrow (four row) constraints. In other words, row groups of at least two rows may be used in embodiments to constrain the knight's order processing method. Assuming r as the number of rows used to constrain the knight's order processing method, the algorithm may be generalized as:

10

15

If not on the bottom row of a row group:

The next block is p columns left, one row down $(-p,+1)$.

Otherwise, at the bottom row of a row group:

The next block is q columns right, $(r-1)$ rows up $(+q,-(r-1))$.

20

[0079] Changing the value of p would affect the value of q , would not affect spacing between adjacent blocks from a row in the pipeline, but would affect spacing between a given block and its other neighbor blocks (e.g., its top-left, top, and top-right neighbors). In particular, note that using the spacing $(-1,+1)$ would result in a block and its diagonal (top-right) neighbor block being concurrently processed at adjacent stages of the pipeline. Thus, a spacing of at least two blocks left may be used so that diagonally adjacent blocks are not concurrently processed at adjacent stages of the block processing pipeline. Changing the value of r would affect the value of q , would affect spacing between adjacent blocks from a row in the pipeline, and would affect spacing between the block and its other neighbor blocks (e.g., its top-left, top, and top-right neighbors).

25

30

[0080] The above algorithm for determining a next block may begin at an initial block. Upon reaching the end of a quadrow that is followed by another quadrow, the algorithm jumps to the first block of the next quadrow and then crosses over between the quadrow and the next quadrow for a few cycles, resulting in the interleaving of some blocks from the end of the quadrow with some blocks from the beginning of the next quadrow. In other words, the knight's order processing method treats the quadrows as if they were arranged end to end. To avoid

35

complications in the algorithm and to maintain consistent spacing of blocks in the pipeline, at least some embodiments may pad the beginning of the first quadrow and the end of the last quadrow with invalid blocks. An invalid block may be defined as a block that is outside the boundary of the frame and that is input to the pipeline but that does not contain valid frame data, and thus is not processed at the stages. The algorithm for determining a next block may thus begin at an initial block, which may be either the first block in the top row of the first quadrow or an invalid block to the left of the first block in the top row of the first quadrow, proceed through all of the quadrows, and at the end of the last quadrow continue until the last block of the last quadrow has been input to the pipeline. There will be bubbles in the pipeline at the beginning and end of the frame, but the spacing of the valid blocks from the frame in the pipeline will remain consistent throughout. In some embodiments, as an alternative to padding the end of the last quadrow of a video frame with invalid blocks, the last quadrow of a video frame may be overlapped with the first row of the next video frame to be processed in the block processing pipeline.

[0081] Figs. 10A and 18B graphically illustrate the knight's order processing method, according to at least some embodiments. For simplicity, these Figures use an example 192x192 pixel frame divided into 144 16x16 pixel blocks, with 12 rows and 12 columns of blocks. However, it is to be noted that the knight's order processing method can be applied to input video frames of any dimensions. In Fig. 10A, an example frame is divided into rows and columns of blocks. The rows of blocks are partitioned into three quadrows including four rows each. The last three rows of the first quadrow are padded on the left with invalid blocks, and the first three rows of the last (third) quadrow are padded on the right with invalid blocks. In this example, the numbers in the blocks represent the order in which the blocks are input to the block processing pipeline according to the knight's order processing method, beginning with block 0 (the first block in the top row of the first quadrow). Block 0 is input to the first stage of the pipeline, and when the first stage is ready for another block, the method proceeds by going two columns left, one row down to get the next block for input (block 1, in Fig. 10A). This pattern is repeated until reaching the bottom of the quadrow. At the bottom of the quadrow, the method goes seven columns right, three rows up to get the next block. This continues until all of the blocks in the frame (as well as all of the invalid blocks shown in Fig. 10A) are input into the pipeline. When the end of a quadrow is reached, if there is another quadrow after the quadrow the input algorithm proceeds to the beginning of the next quadrow. In this example, after block 47 is input, the method proceeds to block 48 (the first block in the top row of the second quadrow). As shown by the dashed arrow from block 47 to the dashed rectangle labeled 48 to the right of block 44, the first block of the top row of the second quadrow (block 48) is treated as being

immediately to the right of the last block of the top row of the first quadrow (block 44), and thus is reached from block 47 by going seven columns right, three columns up. In other words, the knight's order processing method treats the quadrows as if they were arranged end to end, with invalid blocks at each end, as shown in Fig. 10B. Thus, the algorithm for determining a next block remains the same across the entire frame.

[0082] In some embodiments, each row of the first quadrow may be padded with extra invalid blocks, for example with two extra invalid blocks. Instead of beginning with the first block in the top row of the first quadrow as shown in Fig. 10A, input to the pipeline may begin with the first invalid block to the left of the first block in top row of the first quadrow.

[0083] Figs. 11A and 11B are high-level flowcharts of a knight's order processing method for a block processing pipeline, according to at least some embodiments. In Fig. 11A, as indicated at 3100, a next block is determined according to the algorithm for determining a next input block that is implemented by the knight's order processing method. As indicated at 3102, the block is input to the pipeline, for example from a memory via direct memory access (DMA). As shown by 3104, the input process of elements 3100 and 3102 continues as long as there are blocks to be processed. Each block that is input to the pipeline by elements 3100 and 3102 is processed in the pipeline, as indicated at 3106. Each block is initially input to a first stage of the pipeline, processed, output to a second stage, processed, and so on. When a block moves from a stage to a next stage of the pipeline, the stage can begin processing the next block in the pipeline. Thus, the input blocks move through the stages of the pipeline, with each stage processing one block at a time. As indicated at 3108, once a block has been processed by a last stage of the pipeline, the processed block is output, for example to a memory via direct memory access (DMA).

[0084] Fig. 11B is a flowchart of an example algorithm for determining a next input block that that may be implemented by the knight's order processing method, and expands on element 3100 of Fig. 11A. Fig. 11B assumes that the frame is divided into quadrows, and that the algorithm used to determine the next frame is two columns left, one row down (-2,+1) if not on the bottom row of a quadrow, seven columns right, three rows up (+7,-3) if on the bottom row. However, other row groupings and/or spacing algorithms may be used. At 3150, if at the start of the frame, the method gets an initial block as indicated at 3152. If this is not the start of the frame, then at 3154, if this is the last row of the quadrow, the next block is seven columns right, three rows up, as indicated at 3156. If this is not the last row of the quadrow, the next block is two columns left, one row down, as indicated at 3158.

Caching neighbor data

[0085] One or more operations performed at stages of a block processing pipeline may depend on one or more of the neighbor blocks from the previous (or above) row of blocks such as the top neighbor, top-left neighbor, top-right neighbor, and top-right-right neighbor blocks, as well as on the left neighbor, as shown in Fig. 9. The knight's order processing method with quadrow constraints provides locality of neighbor information that may be leveraged to provide local caching of neighbor data at each stage of the pipeline in relatively small local buffers. In at least some embodiments, the local buffers may be implemented using SRAM (static random access memory) technology. However, the local buffers may be implemented using other memory technologies in some embodiments.

[0086] Note that blocks in the first column of a frame do not have a left or top-left neighbor, blocks in the last column do not have a top-right or top-right-right neighbor, and blocks in the next-to-last column do not have a top-right-right neighbor. Thus, for block processing methods that use information from these neighbor positions, the information in the local buffers for these neighbor positions relative to blocks in those columns is not valid and is not used in processing the blocks in those columns in the stages of the pipeline. In addition, there are no rows above the top row of the first quadrow, so the blocks in this row do not have top, top-left, top-right, and top-right-right neighbors.

[0087] In at least some embodiments of a block processing pipeline that implements the knight's order processing method, a first buffer of sufficient size to cache the C most recently processed blocks on the current quadrow may be implemented at each of one or more stages of the pipeline. This buffer may be referred to as the current quadrow buffer, and may, for example, be implemented as a circular FIFO buffer. In at least some embodiments, C may be determined such that the buffer includes an entry corresponding to the top-left neighbor of the current block at the stage according to the algorithm for determining a next block and the row group size used to constrain the knight's order method. The buffer may also include entries corresponding the top-right-right, left, top-right, and top neighbors for the current block according to the algorithm. When processing a block, a stage may access the current quadrow buffer to obtain neighbor information for the block if that block's neighbor information is valid in the current quadrow buffer. Note that some block processing methods may not require top-left neighbor information, and the current quadrow buffer may be smaller in these implementations.

[0088] When a stage completes processing of a block, the block's information is written to the last position in the current quadrow buffer, overwriting the entry at the position of the block's top-left neighbor, thus preparing the buffer for the next block to be processed at the stage. Note that, initially, at the beginning of a frame, there is no information in the current quadrow buffer

as no blocks in the frame have been processed, so no block information will be overwritten in the buffer until the buffer is filled. When the next block is at the stage, the previous block's information in the buffer is the block's top-right-right neighbor information.

[0089] For example, using quadrow boundaries and the algorithm for determining a next block where the next block is two columns left, one row down if not on the bottom row of a quadrow, $C = 13$ would be sufficient to include the top-left neighbor of the current block, as the spacing between the current block and its top-left neighbor is 13. Fig. 12 shows a portion of a quadrow as processed in a pipeline according to the knight's order processing method that may be cached in the current quadrow buffer, according to at least some embodiments. Block 19 represents a current block at a stage. The shaded blocks represent the 13 most recently processed blocks by the stage. Note that the farthest block from block 19 in time is its top-left neighbor (block 6), and the nearest block in time is its top-right-right neighbor (block 9).

[0090] For the blocks in the top row of a quadrow, information for neighbors in the row above is not in the current quadrow buffer. There are no rows above the top row of the first quadrow, and for all other quadrows the row above the top row is the bottom row of the previous quadrow. Thus, the current quadrow buffer includes the left neighbor information for all blocks in the top row of a quadrow (except for the first block, which has no left neighbor), but does not include the top-left, top, top-right, and top-right-right neighbor information for the blocks in the top row of the quadrow. To provide this neighbor information for blocks on the top rows of the quadrows, a second buffer of sufficient size to hold information for the required neighbor blocks from the last row of the previous quadrow may be implemented at one or more stages of the pipeline. This buffer may be referred to as the previous quadrow buffer, and may, for example, be implemented as a circular FIFO buffer. The number of entries in the previous quadrow buffer, as well as the particular neighbor blocks that are cached in the previous quadrow buffer, may be dependent on the requirements of the particular block processing method that is implemented by the block processing pipeline. In at least some embodiments, when processing a quadrow according to the knight's order processing method, information for each block on the bottom row of the quadrow may be written to an external memory, for example when the block is at a last stage of the pipeline. For each block in the top row of a quadrow, neighbor (e.g., top-right-right neighbor) data may be read from the external memory, for example at a first stage of the pipeline. This neighbor information may be passed down the pipeline to the other stages along with the corresponding block from the top row.

[0091] Fig. 13 graphically illustrates blocks in a current quadrow being processed according to the knight's order processing method, as well as neighbor blocks in the last row of the previous quadrow, according to at least some embodiments. Blocks A , $A+4$, $A+8$, and $A+12$ were

processed on the previous quadrow according to the knight's order processing method. Block A was processed first, block $A+4$ was processed four cycles later, and so on. Block B represents a block on the current quadrow that is currently at a particular stage of the pipeline. Blocks $B-1$ (B minus 1) through $B-13$ (B minus 13) represent the thirteen blocks that were most recently processed at the stage in the current quadrow. Information from these blocks may be presently cached in the stage's current quadrow buffer, with $B-1$ as the most recent entry and $B-13$ as the oldest entry. $B-4$ is current block B 's left neighbor. However, block B 's top-left (block A), top (block $A+4$), top-right (block $A+8$), and top-right-right (block $A+12$) neighbors are on the bottom row of the previous quadrow, and are not included in the current quadrow buffer for block B . In at least some embodiments, to provide neighbor information for blocks on the top row of the current quadrow (e.g., top-left, top, top-right, and top-right-right neighbor information), a previous quadrow buffer may be implemented at each of one or more stages of the pipeline. When processing a quadrow, information for each block on the bottom row of the quadrow is written to a neighbor data structure in external memory, for example by a last stage of the pipeline. When processing blocks from the top row of a next quadrow, information for neighbor blocks in the bottom row of the previous quadrow is read from the external memory, for example by a first stage of the pipeline, and passed down the pipeline to other stages with the top row blocks. In at least some embodiments, information for the top-right-right neighbor block of a block in the top row is read from the external memory. In at least some embodiments, the previous quadrow buffer is a circular buffer, and an oldest entry in the previous quadrow buffer is replaced with the neighbor information that is read from the external memory. In various embodiments, the external memory to which blocks in the bottom row are written and from which neighbor block information is read may be a memory of the pipeline component that is external to the last stage, a memory of a video encoder that implements the pipeline, or a memory external to the video encoder. In some embodiments, however, the memory may be a local memory of the last stage of the pipeline. At least some embodiments may include an interlock mechanism to control the reads and writes to the external memory between rows to avoid overwriting the data in external memory.

[0092] Fig. 14 is a flowchart of a method for processing blocks in a block processing pipeline in which neighbor data is cached in local buffers at the stages of the pipeline, according to at least some embodiments. For example, the method of Fig. 14 may be used at element 3106 of Fig. 11A to process blocks input to the pipeline according to the knight's order processing method as shown at elements 3100, 3102, and 3104 of Fig. 11A. In Fig. 14, a block is input to the pipeline. At 4200, at a first stage of the pipeline, if the block is on the top row of a quadrow, then neighbor data for the block may be read from external memory (for example, via DMA) into

a previous quadrow buffer as indicated at 4202. In at least some embodiments, the neighbor data corresponds to the top-right-right neighbor of the current block on the bottom row of the previous quadrow. As indicated at 4204, the block is then processed at the current stage. If an operation at the stage requires neighbor information to process the block, the stage may use the neighbor information in the current quadrow buffer and/or in the previous quadrow buffer to perform the operation. If the block is on the top row of a quadrow, then at least some of the neighbor information is obtained from the previous quadrow buffer; otherwise, neighbor information may be obtained from the current quadrow buffer. As indicated at 4206, information about the current block may be written to the current quadrow buffer at the stage for use on subsequent blocks. The information may overwrite an oldest entry in the current quadrow buffer.

[0093] At 4208, if there are more stages, then the block may be sent to a next stage, as indicated at 4210. At 4212, neighbor information from the previous quadrow buffer may also be sent to the next stage. In at least some embodiments, this neighbor information is only sent to the next stage if the current block is on the top row of a quadrow. Elements 4204 through 4212 may be repeated until the block reaches and is processed by a last stage of the pipeline. At 4208, if there are no more stages, then processing of the block in the pipeline is done. At 4214, if the block is on the bottom row of a quadrow, then information for the block is written to an external memory (for example, via DMA) to be read as neighbor data for blocks in the top row of a next quadrow. In addition, all of the processed valid blocks are output as shown by element 3108 of Fig. 11A.

Example pipeline units

[0094] Figs. 15A through 15C are block diagrams of example pipeline processing units that may be used at the stages of a block processing pipeline that implements one or more of the block processing methods and apparatus as described herein, according to at least some embodiments. For example, one or more of pipeline units 5000A and/or 5000B as shown in Figs. 15A and 15B may be used at each stage of the example block processing pipeline shown in Fig. 16. Note that Figs. 15A through 15C are not intended to be limiting; a pipeline processing unit may include more or fewer components and features than those shown in the Figures.

[0095] As shown in Fig. 15A, a pipeline unit 5000A may include at least a memory 5010 and a unit core 5020. Unit core 5020 may be a component (e.g., a circuit) that is configured to perform a particular operation on or for a block, or a portion of a block, at a particular stage of the block processing pipeline. Memory 5010 may, for example, be a double-buffered memory that allows the unit core 5020 to read and process data for a block from the memory 5010 while data for a next block is being written to the memory 5010 from a previous pipeline unit.

[0096] As shown in Fig.15B, a pipeline unit 5000B, in addition to a memory 5010 and unit core 5020 as shown in Fig. 15A, may also include a processor 5030. Processor 5030 may, for example, be a mobile or M-class processor. The processors 5030 in pipeline units 5000B of a block processing pipeline may, for example, be used to control the block processing pipeline at block boundaries. The processors 5030 in pipeline units 5000B may be configurable, for example with low-level firmware microcode, to allow flexibility in algorithms that are implemented by the block processing pipeline for various applications. In at least some embodiments, a processor 5030 of a pipeline unit 5000B in the pipeline may be configured to receive data from a processor 5030 of a previous (upstream) pipeline unit 5000B and send data to a processor 5030 of a subsequent (downstream) pipeline unit 5000B. In addition, a processor 5030 of a pipeline unit 5000B at a last stage of the pipeline may be configured to send feedback data to a processor 5030 of a pipeline unit 5000B at a first stage of the pipeline.

[0097] As shown in Figs. 15A and 15B, a pipeline unit 5000A or 5000B may be configured to access external memory, for example according to direct memory access (DMA). In addition, a pipeline unit 5000A or 5000B may be configured to pass information back to one or more previous (upstream) stages of the pipeline and/or to receive information passed back from one or more subsequent (downstream) stages of the pipeline. In addition, a pipeline unit 5000A or 5000B may be configured to pass information forward to one or more subsequent (downstream) stages of the pipeline and/or to receive information passed forward from one or more previous (upstream) stages of the pipeline.

[0098] As shown in Fig. 15C, two or more units 5000A as shown in Fig. 15A may be grouped together and configured to perform an operation in the pipeline. A single processor 5030 may be used to control and/or configure the pipeline units 5000A.

Example block processing pipeline

[0099] Fig. 16 is a high-level block diagram of general operations in an example block processing method 6000 for H.264 encoding that may be implemented in stages by a block processing pipeline that may implement one or more of the block processing methods and apparatus as described herein, according to at least some embodiments. A block processing pipeline that implements the block processing method 6000 may, for example, be implemented as a component of an H.264 video encoder apparatus that is configured to convert input video frames from an input format into H.264 / Advanced Video Coding (AVC) format as described in the H.264/AVC standard. The H.264/AVC standard is published by ITU-T in a document titled "ITU-T Recommendation H.264: Advanced video coding for generic audiovisual services", which may be referred to as the H.264 Recommendation. An example input video format is

1080p (1920x1080 pixels, 2.1 megapixels) encoded in YCbCr color space. However, other input video formats may be encoded into H.264 using embodiments of the pipeline in a video encoder apparatus.

[00100] The video encoder apparatus may, for example, be implemented as an integrated circuit (IC) or as a subsystem on an IC such as a system-on-a-chip (SOC). In at least some embodiments, the video encoder apparatus may include at least a pipeline component, a processor component (e.g., a low-power multicore processor), and a bus subsystem or fabric that interconnects the functional components of the apparatus. The processor component of the video encoder apparatus may, for example, perform frame-level control of the pipeline such as rate control, perform pipeline configuration, and interface with application software via a driver. The pipeline component may implement multiple processing stages each configured to perform a portion or all of one or more of the operations as shown in Fig. 16, each stage including one or more processing units. At least one of the processing units in the pipeline may include a processor component (e.g., an M-class processor) that may, for example, configure parameters of the processing unit at the respective stage at the macroblock level. The video encoder apparatus may include other functional components or units such as memory components, as well as external interfaces to, for example, one or more video input sources and external memory. Example video input sources to the video encoder apparatus may include one or more of, but are not limited to, a video camera for raw video input processing, a decoder apparatus for re-encoding/transcoding, a flash or other memory, and a JPEG decoder. An example video encoder apparatus is illustrated in Fig. 15. An example SOC that includes a video encoder apparatus is illustrated in Fig. 16. While embodiments are generally described in relation to hardware implementations of a block processing pipeline that implements the block processing method 6000 with knight's order processing, note that the block processing method 6000 with knight's order processing may be implemented by a block processing pipeline implemented in software.

[00101] A pipeline that implements the method 6000 as shown in Fig. 16 may process 16x16 pixel macroblocks from input video frames according to the H.264 standard, each macroblock including two or more blocks or partitions that may be processed separately at stages of the pipeline. The input video frames may, for example, be encoded in YCbCr color space; each macroblock may be composed of separate blocks of chroma and luma elements that may be processed separately at the stages of the pipeline. A pipeline that implements the block processing method 6000 may receive input macroblocks from and output processed macroblocks to a memory. The memory may include memory of the video encoder apparatus and/or memory external to the video encoder apparatus. In at least some embodiments, the memory may be accessed by the pipeline as necessary, for example via direct memory access (DMA). In at least

some embodiments, the memory may be implemented as a multi-level memory with a cache memory implemented between the pipeline and an external memory. For example, in some implementations, one or more quadrows may be read from an external memory and cached to the cache memory for access by the pipeline to reduce the number of reads to an external memory.

5 [00102] The general operations of the example H.264 video encoder method 6000 as shown in Fig. 16 that may be performed in stages by a pipeline, as well as general data flow through the pipeline, are briefly described below. Each of the general operations of the method 6000 may be implemented by one or more pipeline units at one or more stages of the pipeline. Example pipeline units are illustrated in Figs. 13A through 13C. Also note that each general operation
10 shown in Fig. 16 may be subdivided into two or more operations that may be implemented by pipeline units at one, two, or more stages of the pipeline. However, two or more of the operations shown in Fig. 16 may be performed at the same stage of the pipeline. Each stage in the pipeline processes one macroblock at a time, and thus two or more of the operations may simultaneously operate on the same macroblock that is currently at the respective stage. Note
15 that a pipeline may perform more, fewer, or other operations than those shown in Fig. 16 and described below.

Macroblock input

[00103] In at least some embodiments, macroblock input 6002 may be performed by an initial
20 stage of the pipeline. In at least some embodiments, macroblock input 6002 receives luma and chroma pixels from a memory, for example via DMA, computes statistics on input pixels that are used by firmware in downstream stages of the pipeline, and buffers input macroblocks to enable firmware look ahead. The input macroblock pixel data and corresponding statistics are buffered and sent to one or more downstream stages of the pipeline that implement intra-frame and inter-
25 frame estimation 6010 operations. In at least some embodiments, an input buffer of up to 16 macroblocks is maintained for input pixels and statistics. The macroblock pixel data and corresponding statistics may be input to downstream stages of the pipeline according to a knight's order input algorithm as previously described in relation to Figs. 3 through 8B.

[00104] In at least some embodiments, macroblock input 6002 reads neighbor data from the
30 bottom row of a previous quadrow from memory at quadrow boundaries and passes the neighbor data to at least one downstream stage.

Intra-frame and inter-frame estimation

[00105] Intra-frame and inter-frame estimation 6010 operations may determine blocks of
35 previously encoded pixels to be used in encoding macroblocks input to the pipeline. In H.264

video encoding, each macroblock can be encoded using blocks of pixels that are already encoded within the current frame. The process of determining these blocks may be referred to as intra-frame estimation, or simply intra-estimation. However, macroblocks may also be encoded using blocks of pixels from one or more previously encoded frames (referred to as reference frames).

5 The process of finding matching pixel blocks in reference frames may be referred to as inter-frame estimation, or more generally as motion estimation. Intra-frame and inter-frame estimation 6010 operations may be subdivided into two or more sub-operations that may be performed at one, two, or more stages of the pipeline, with one or more components or pipeline units at each stage configured to perform a particular sub-operation.

10 **[00106]** In at least some embodiments, macroblock input 6002 reads neighbor data from the bottom row of a previous quadrow from memory at quadrow boundaries and passes the neighbor data to intra-frame and inter-frame estimation 6010, for example to an intra-frame estimation component. In addition, motion compensation and reconstruction 6030, for example a luma reconstruction component, may pass neighbor data as feedback to intra-frame and inter-frame 15 estimation 6010, for example to the intra-frame estimation component.

Motion estimation

[00107] In at least some embodiments, to perform motion estimation, the pipeline may include one instance of a motion estimation engine for each reference frame to be searched. Each motion 20 estimation engine searches only one reference frame. In at least some embodiments, each motion estimation engine may include a low resolution motion estimation component, a full pixel motion estimation component, and a subpixel motion estimation component. In at least some embodiments, the three components of each of the motion estimation engines may be implemented at different stages of the pipeline. In at least some embodiments, each motion 25 estimation engine may also include a memory component that reads and stores reference frame data from a memory as needed. In at least some embodiments, a single instance of a processor manages all instances of the motion estimation engine. In at least some embodiments, the processor may determine one or more candidates using predicted and co-located motion vectors and input the candidates to the full pixel motion estimation components of the motion estimation 30 engines.

[00108] In at least some embodiments, the low resolution motion estimation component of each motion estimation engine performs an exhaustive search on a scaled-down, low resolution version of a respective reference frame to generate candidates. In at least some embodiments, the full pixel motion estimation component performs a search on full size pixels using candidates 35 from the low resolution motion estimation component. In at least some embodiments, the

subpixel motion estimation component performs a search on half and quarter pixels using best candidates received from the full pixel motion estimation component. In some embodiments, full pixel motion estimation and subpixel motion estimation may be disabled based on results of a direct mode estimation performed at an upstream stage of the pipeline. In at least some
5 embodiments, each motion estimation engine outputs results data to mode decision 6020.

[00109] In at least some embodiments, motion estimation may also include a direct mode estimation component that receives co-located and spatial motion vector data and computes a direct/skip mode cost, which it provides to mode decision 6020. Based on the results, the direct mode estimation component may disable full pixel motion estimation and subpixel motion
10 estimation.

Intra estimation

[00110] In at least some embodiments, an intra estimation component of the pipeline performs intra mode selection to determine blocks of pixels already encoded within the current frame that
15 may be used in encoding a current macroblock. In at least some embodiments, the intra estimation component performs intra mode selection only for luma. In these embodiments, Chroma intra estimation is performed by a chroma reconstruction component at a downstream stage of the pipeline. In at least some embodiments, the intra estimation component may perform intra estimation independently for each of two or more blocks or partitions (e.g., 4x4, 8x8, 4x8,
20 8x4, 16x8, and/or 8x16 blocks) in a macroblock. For each block, prediction pixels are first extracted from neighbor blocks (neighbor blocks can be outside the current macroblock in the frame or within the current macroblock). For each prediction mode in the current block, the cost of the current mode is evaluated by creating a prediction block from neighbor pixels, computing a mode cost, and comparing the mode cost to a minimum cost for that block. Once all prediction
25 modes are evaluated and the best mode is determined, reconstruction may be performed for the best mode so that reconstructed pixels can be used to predict future blocks within the macroblock. The intra estimation component may pass best intra mode information to mode decision 6020.

[00111] In at least some embodiments, macroblock input 6002 reads neighbor data from the
30 bottom row of a previous quadrow from memory at quadrow boundaries and passes the neighbor data to the intra estimation component. In at least some embodiments, at least one downstream stage (e.g., a luma reconstruction component at a downstream stage) may pass neighbor data back to the intra estimation component.

Mode decision

[00112] In at least some embodiments, mode decision 6020 may be implemented by a mode decision component at a stage of the pipeline that is downstream of the stage(s) that implement intra-frame and inter-frame estimation 6010 operations. However, in some embodiments, mode decision 6020 operations may be subdivided into two or more sub-operations that may be performed at one, two, or more stages of the pipeline, with one or more components or pipeline units at each stage configured to perform a particular sub-operation. In at least some embodiments, the mode decision 6020 component receives the best intra mode from intra estimation, direct/skip mode cost from direct mode estimation, and motion vector candidates from the motion estimation engines. In at least some embodiments, the mode decision component computes additional costs for bi-directional modes and determines the best macroblock type, including macroblock partitions, sub-partitions, prediction direction and reference frame indices. In at least some embodiments, the mode decision 6020 component also performs all motion vector prediction. The motion vector prediction results may be used when estimating motion vector rate during mode decision. In at least some embodiments, the motion vector prediction results may also be fed back from the mode decision 6020 component to motion estimation, for example for use in direct mode estimation and motion vector rate estimation.

Motion compensation and reconstruction

[00113] In at least some embodiments, motion compensation and reconstruction 6030 operations may be subdivided into two or more sub-operations that may be performed at one, two, or more stages of the pipeline, with one or more components or pipeline units at each stage configured to perform a particular sub-operation. For example, in some embodiments, motion compensation and reconstruction 6030 may be subdivided into luma motion compensation and reconstruction and chroma motion compensation and reconstruction. In at least some embodiments, each of these sub-operations of motion compensation and reconstruction 6030 may be performed by one or more components or pipeline units at one or more stages of the pipeline.

Luma motion compensation and reconstruction

[00114] In at least some embodiments, a luma motion compensation component of the pipeline receives the best mode and corresponding motion vectors from mode decision 6020. As previously noted, each motion estimation engine may include a memory component that reads and stores reference frame data from a memory. If the best mode is inter-predicted, the luma motion compensation component requests reference frame macroblocks from the motion

estimation engine corresponding to the motion vectors. The motion estimation engine returns subpixel interpolated 4x4 or 8x8 blocks depending on the request size. The luma motion compensation component then combines the blocks into prediction macroblocks. The luma motion compensation component then applies a weighted prediction to the prediction macroblocks to create the final macroblock predictor that is then passed to the luma reconstruction component.

[00115] In at least some embodiments, a luma reconstruction component of the pipeline performs macroblock reconstruction for luma, including intra prediction (in at least some embodiments, the luma motion compensation component performs inter prediction), forward transform and quantization (FTQ), and inverse transform and quantization (ITQ).

[00116] In at least some embodiments, based on the best mode from mode decision 6020, either an inter prediction macroblock is passed from the luma motion compensation component or intra prediction is performed by the luma reconstruction component to generate a prediction block. In intra mode, the prediction is performed in block (scan) order since reconstructed pixels from neighbor blocks are needed for prediction of future blocks. The input block is subtracted from the prediction block to generate a residual block. This residual pixel data is transformed and quantized by an FTQ technique implemented by the luma reconstruction component. The coefficient data is sent to an ITQ technique implemented by the luma reconstruction component, and may also be sent downstream to CAVLC encoding. The ITQ technique generates a reconstructed residual pixel block. The prediction block is added to the residual block to generate the reconstructed block. Reconstructed pixels may be passed downstream to a deblocking filter. In at least some embodiments, reconstructed pixels may also be passed back to an intra-frame estimation component of intra-frame and inter-frame estimation 6010 for prediction of future blocks inside the current macroblock.

Chroma motion compensation and reconstruction

[00117] In at least some embodiments, chroma reconstruction is performed in two stages. In the first stage, chroma reference blocks needed for inter prediction are read from memory based on input macroblock type, motion vectors, and reference frame index. Subpixel interpolation and weighted prediction is then applied to generate a prediction macroblock. In the second stage, chroma intra prediction and chroma intra/inter FTQ/ITQ is performed. This allows one additional pipeline stage to load chroma prediction pixel data. Since chroma pixels are not searched by motion estimation, the chroma prediction data is read from external memory and may have large latency. In at least some embodiments, a chroma motion compensation component performs the first stage, while a chroma reconstruction component performs the second stage.

[00118] In at least some embodiments, the chroma motion compensation component generates a prediction block including subpixel interpolation for Cb and Cr chroma blocks; the size is based on the partition size and chroma formats. A full size chroma block is 8x8, 8x16, or 16x16 pixels for chroma formats 4:2:0, 4:2:2 and 4:4:4, respectively. In at least some embodiments, the chroma motion compensation component may prefetch and cache chroma prediction pixels from an external (to the pipeline) memory. In at least some embodiments, reference data may be read based on mode decision 6020 results. The chroma motion compensation component performs subpixel interpolation to generate a prediction block. Mode decision 6020 provides the macroblock type and sub-types, reference frame index per partition, and corresponding motion vectors. The prediction is output to the chroma reconstruction component.

[00119] In at least some embodiments, the chroma reconstruction component performs chroma prediction, chroma intra estimation and chroma reconstruction for inter and intra modes. For chroma formats 4:2:0 and 4:2:2, intra chroma estimation and prediction is performed. In at least some embodiments, chroma intra estimation is performed at this stage rather than at intra-frame and inter-frame estimation 6010 so that reconstructed pixels can be used during the estimation process. In at least some embodiments, if the best mode is in intra, intra chroma estimation may be performed. based on the best intra chroma mode, and intra prediction may be performed using one of four intra chroma modes. For inter macroblocks, inter chroma prediction pixels are received from chroma motion compensation. For chroma format 4:4:4, the luma intra prediction modes are used to generate the chroma block prediction, and inter chroma prediction is performed in the same manner as for luma. Therefore, chroma reconstruction conceptually includes 4:2:0 and 4:2:2 chroma reconstruction and luma reconstruction used to reconstruct chroma in 4:4:4 chroma format.

CAVLC encode and deblocking

[00120] In at least some embodiments, CAVLC encoding and deblocking may be performed by one or more components at a last stage of the pipeline. In at least some embodiments, a deblocking filter component of the pipeline receives reconstructed luma and chroma pixels from the chroma reconstruction component and performs deblocking filtering according to the H.264 Recommendation. Results may be output to a memory.

[00121] In at least some embodiments, a CAVLC encode component of the pipeline receives at least luma and chroma quantized coefficients, neighbor data, and chroma reconstruction results from the chroma reconstruction component and generates a CAVLC (context-adaptive variable-length coding) encoded output stream to a memory.

[00122] In at least some embodiments, the deblocking filter component and the CAVLC encode component write neighbor data for the bottom row of a quadrow to a memory at quadrow boundaries. For the top row of a next quadrow, macroblock input 6002 may then read this neighbor data from the memory at quadrow boundaries and pass the neighbor data to at least one downstream stage of the pipeline.

Transcoder

[00123] In at least some embodiments, a transcoding operation may be performed by a transcoder 6050. The transcoder may be implemented as a functional component of the pipeline or as a functional component that is external to the pipeline. In at least some embodiments, the transcoder 6050 may perform a memory-to-memory conversion of a CAVLC (context-adaptive variable-length coding) encoded stream output by the pipeline to a CABAC (context-adaptive binary arithmetic coding) encoded stream.

[00124] In at least some embodiments, the pipeline may encode in an order other than scan order, for example knight's order as previously described herein. However, ultimately, the H.264 video encoder's encoded bit stream should be transmitted in conventional macroblock scan order. In at least some embodiments, re-ordering the macroblock output from knight's order to scan order is accomplished by the CAVLC encode component writing encoded data to four different output buffers, each output buffer corresponding to a macroblock row. At the end of a quadrow, each row buffer will contain a scan order stream of encoded macroblocks for a respective row. Transcoder 6050 handles stitching the start and end of each row to generate a continuous stream at macroblock row boundaries. In at least some embodiments, the pipeline may embed metadata in the CAVLC output stream to facilitate stitching of the rows by the transcoder 6050.

Example video encoder apparatus

[00125] Fig. 17 is a block diagram of an example video encoder apparatus 7000, according to at least some embodiments. The video encoder apparatus 7000 may, for example, be implemented as an integrated circuit (IC) or as a subsystem on an IC such as a system-on-a-chip (SOC). In at least some embodiments, the video encoder apparatus 7000 may include a pipeline 7040 component, a processor 7010 component (e.g., a low-power multicore processor), a memory management unit (MMU) 7020, DMA 7030, and an interconnect 7050 such as a bus subsystem or fabric that interconnects the functional components of the apparatus. The processor 7010 component of the video encoder apparatus 7000 may, for example, perform frame-level control of the pipeline 7040 such as rate control, perform pipeline 7040 configuration including

configuration of individual pipeline units within the pipeline 7040, and interface with application software via a driver, for example for video encoder 7000 configuration. The MMU 7020 may serve as an interface to external memory, for example for streaming video input and/or output. Pipeline 7040 component may access memory through MMU 7020 via DMA 7030. In some
5 embodiments, the video encoder apparatus 7000 may include other functional components or units not shown in Fig. 17, or fewer functional components than those shown in Fig. 17. An example block processing method that may be implemented by pipeline 7040 component is shown in Fig. 16. An example a system-on-a-chip (SOC) that may include at least one video encoder apparatus 7000 is illustrated in Fig. 18.

Example system on a chip (SOC)

[00126] Turning now to Fig. 18, a block diagram of one embodiment of a system-on-a-chip (SOC) 8000 that may include at least one instance of a video encoder apparatus including a block processing pipeline that may implement one or more of the block processing methods and
15 apparatus as illustrated in Figs. 3 through 17. SOC 8000 is shown coupled to a memory 8800. As implied by the name, the components of the SOC 8000 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. However, the SOC 8000 will be used as an example herein. In the illustrated embodiment, the components of the SOC 8000
20 include a central processing unit (CPU) complex 8020, on-chip peripheral components 8040A-8040B (more briefly, "peripherals"), a memory controller (MC) 8030, a video encoder 7000 (which may itself be considered a peripheral component), and a communication fabric 8010. The components 8020, 8030, 8040A-8040B, and 7000 may all be coupled to the communication fabric 8010. The memory controller 8030 may be coupled to the memory 8800 during use, and
25 the peripheral 8040B may be coupled to an external interface 8900 during use. In the illustrated embodiment, the CPU complex 8020 includes one or more processors (P) 8024 and a level two (L2) cache 8022.

[00127] The peripherals 8040A-8040B may be any set of additional hardware functionality included in the SOC 8000. For example, the peripherals 8040A-8040B may include video
30 peripherals such as an image signal processor configured to process image capture data from a camera or other image sensor, display controllers configured to display video data on one or more display devices, graphics processing units (GPUs), video encoder/decoders, scalers, rotators, blenders, etc. The peripherals may include audio peripherals such as microphones, speakers, interfaces to microphones and speakers, audio processors, digital signal processors,
35 mixers, etc. The peripherals may include peripheral interface controllers for various interfaces

8900 external to the SOC 8000 (e.g. the peripheral 8040B) including interfaces such as Universal Serial Bus (USB), peripheral component interconnect (PCI) including PCI Express (PCIe), serial and parallel ports, etc. The peripherals may include networking peripherals such as media access controllers (MACs). Any set of hardware may be included.

5 [00128] More particularly in Fig. 18, SOC 8000 may include at least one instance of a video encoder 7000 component, for example a video encoder 7000 as illustrated in Fig. 17 that includes a block processing pipeline 7040 component that implements a block processing method 6000 as illustrated in Fig. 16. Video encoder 7000 may be an H.264 video encoder apparatus that may be configured to convert input video frames from an input format into H.264 / Advanced Video
10 Coding (AVC) format as described in the H.264/AVC standard. The block processing pipeline 7040 may implement one or more of the block processing methods and apparatus as described herein in relation to Figs. 3 through 16.

[00129] The CPU complex 8020 may include one or more CPU processors 8024 that serve as the CPU of the SOC 8000. The CPU of the system includes the processor(s) that execute the
15 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 8024 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 8024 may also
20 be referred to as application processors. The CPU complex 8020 may further include other hardware such as the L2 cache 8022 and/or and interface to the other components of the system (e.g. an interface to the communication fabric 8010). 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
25 processors in response to executing the instructions may generally be stored in the memory 8800, 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 8000) or other levels of integration. Processors may further encompass discrete microprocessors, processor cores and/or microprocessors integrated
30 into multichip module implementations, processors implemented as multiple integrated circuits, etc.

[00130] The memory controller 8030 may generally include the circuitry for receiving memory operations from the other components of the SOC 8000 and for accessing the memory 8800 to complete the memory operations. The memory controller 8030 may be configured to
35 access any type of memory 8800. For example, the memory 8800 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 8030 may include queues for memory operations, for ordering (and potentially reordering) the operations and presenting the operations to the memory 8800. The memory controller 8030 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 8030 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 8800 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 8022 or caches in the processors 8024, which serve only certain components. Additionally, in some embodiments, a system cache need not be located within the memory controller 8030.

[00131] In an embodiment, the memory 8800 may be packaged with the SOC 8000 in a chip-on-chip or package-on-package configuration. A multichip module configuration of the SOC 8000 and the memory 8800 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 8800 unencrypted, whereas the protected data may be encrypted for exchange between the SOC 8000 and external endpoints.

[00132] The communication fabric 8010 may be any communication interconnect and protocol for communicating among the components of the SOC 8000. The communication fabric 8010 may be bus-based, including shared bus configurations, cross bar configurations, and hierarchical buses with bridges. The communication fabric 8010 may also be packet-based, and may be hierarchical with bridges, cross bar, point-to-point, or other interconnects.

[00133] It is noted that the number of components of the SOC 8000 (and the number of subcomponents for those shown in Fig. 18, such as within the CPU complex 8020) may vary from embodiment to embodiment. There may be more or fewer of each component/subcomponent than the number shown in Fig. 18.

Example system

[00134] Fig. 19 a block diagram of one embodiment of a system 9000. In the illustrated embodiment, the system 9000 includes at least one instance of the SOC 8000 coupled to one or more external peripherals 9020 and the external memory 8800. A power management unit

(PMU) 9010 is provided which supplies the supply voltages to the SOC 8000 as well as one or more supply voltages to the memory 8800 and/or the peripherals 9020. In some embodiments, more than one instance of the SOC 8000 may be included (and more than one memory 8800 may be included as well).

5 **[00135]** The peripherals 9020 may include any desired circuitry, depending on the type of system 9000. For example, in one embodiment, the system 9000 may be a mobile device (e.g. personal digital assistant (PDA), smart phone, etc.) and the peripherals 9020 may include devices for various types of wireless communication, such as wifi, Bluetooth, cellular, global positioning system, etc. The peripherals 9020 may also include additional storage, including RAM storage,
10 solid state storage, or disk storage. The peripherals 9020 may include user interface devices such as a display screen, including touch display screens or multitouch display screens, keyboard or other input devices, microphones, speakers, etc. In other embodiments, the system 9000 may be any type of computing system (e.g. desktop personal computer, laptop, workstation, net top etc.).

15 **[00136]** The external memory 8800 may include any type of memory. For example, the external memory 8800 may be SRAM, dynamic RAM (DRAM) such as synchronous DRAM (SDRAM), double data rate (DDR, DDR2, DDR3, etc.) SDRAM, RAMBUS DRAM, low power versions of the DDR DRAM (e.g. LPDDR, mDDR, etc.), etc. The external memory 8800 may include one or more memory modules to which the memory devices are mounted, such as single inline memory modules (SIMMs), dual inline memory modules (DIMMs), etc. Alternatively, the
20 external memory 8800 may include one or more memory devices that are mounted on the SOC 8000 in a chip-on-chip or package-on-package implementation.

25 **[00137]** The methods described herein may be implemented in software, hardware, or a combination thereof, in different embodiments. In addition, the order of the blocks of the methods may be changed, and various elements may be added, reordered, combined, omitted, modified, etc. Various modifications and changes may be made as would be obvious to a person skilled in the art having the benefit of this disclosure. The various embodiments described herein are meant to be illustrative and not limiting. Many variations, modifications, additions, and improvements are possible. Accordingly, plural instances may be provided for components described herein as a single instance. Boundaries between various components, operations and
30 data stores are somewhat arbitrary, and particular 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 example configurations may be implemented as a combined structure or component. These and other variations, modifications, additions, and improvements may fall
35 within the scope of embodiments as defined in the claims that follow.

WHAT IS CLAIMED IS:

1. An apparatus, comprising:

a block processing pipeline comprising a plurality of stages, wherein one or more of the plurality of stages of the block processing pipeline each comprises a component of a software pipeline and a component of a hardware pipeline, wherein the software and hardware pipelines are configured to process blocks of pixels from a frame in parallel;

wherein the software pipeline component at each of the one or more stages is configured to iteratively determine configurations for processing the blocks at the hardware pipeline component of the stage according to obtained information for the blocks;

wherein the hardware pipeline component at each of the one or more stages is configured to iteratively process the blocks according to the configurations for processing the blocks that are determined by the software pipeline component at the stage; and

wherein the hardware pipeline component at each of the one or more stages processes a current block at the stage while the software pipeline component at the stage is determining a configuration for an upcoming block to be processed by the hardware pipeline component at the stage.

2. The apparatus as recited in claim 1, wherein the software pipeline component at each of the one or more stages is further configured to receive at least a portion of the information for the blocks from an upstream stage of the block processing pipeline.

3. The apparatus as recited in claim 1 or 2, wherein the information for the blocks includes information from blocks that were previously processed by the hardware pipeline component at the stage.

4. The apparatus as recited in any of claims 1 – 3, wherein at least one of the software pipeline components is further configured to output the information for the blocks to a downstream stage of the block processing pipeline.

5. The apparatus as recited in any of claims 1 – 4, wherein the hardware pipeline component at each of the one or more stages is further configured to:

receive the blocks to be processed from a previous stage of the block processing pipeline;

and

output the processed blocks to a next stage of the block processing pipeline or to an external memory.

6. The apparatus as recited in any of claims 1 – 5, wherein the block processing
5 pipeline includes an initial stage configured to:
buffer blocks for input to the hardware pipeline;
calculate block statistics for one or more of the buffered blocks; and
input the block statistics to a software pipeline component at a downstream stage of the
block processing pipeline as information for the blocks.

10

7. The apparatus as recited in claim any of claims 1 – 6, wherein the software
pipeline component at each stage is configured to, for each block to be processed by the
hardware pipeline component at the stage:

15

determine a particular configuration for processing the block at the hardware pipeline
component of the stage according to obtained information for the block;
write the configuration to a configuration memory; and
signal the hardware pipeline component at the stage that the configuration is ready in the
configuration memory.

20

8. The apparatus as recited in claim 7, wherein the hardware pipeline component at
the stage is configured to:

25

detect the signal from the software pipeline component at the stage;
access the configuration from the configuration memory;
signal the software pipeline component that the configuration memory is clear; and
process the block according to the configuration from the configuration memory.

30

9. The apparatus as recited in claim 8, wherein the software pipeline component at
the stage is further configured to write a configuration for a next block to a first part of the
configuration memory while the hardware pipeline component accesses a configuration for a
current block from a second part of the configuration memory.

35

10. A method, comprising:
obtaining, by a software pipeline component at a stage of a block processing pipeline,
information for a block to be processed by a hardware pipeline component at the
stage, wherein the block processing pipeline processes blocks of pixels from a

frame in parallel software and hardware pipelines, wherein one or more of a plurality of stages of the block processing pipeline each comprises a software pipeline component and a hardware pipeline component;

determining, by the software pipeline component, a configuration for processing the block at the hardware pipeline component according to the obtained information for the block;

signaling the hardware pipeline component that the configuration for the block is ready; and

in response to said signaling, processing the block at the hardware pipeline component according to the determined configuration for the block.

wherein the hardware pipeline component processes the block while the software pipeline component is determining a configuration for an upcoming block to be processed by the hardware pipeline component.

11. The method as recited in claim 10, wherein the information for the block includes block information received from an upstream stage of the block processing pipeline and information from another block that was previously processed by the hardware pipeline component at the stage.

12. The method as recited in claim 10 or 11, further comprising:

writing, by the software pipeline component, the configuration for the block to a configuration memory; and

reading, by the hardware pipeline component, the configuration for the block from the configuration memory in response to said signaling.

13. A device, comprising:

a memory; and

an apparatus configured to process video frames and to store the processed video frames as frame data to the memory, the apparatus comprising a block processing pipeline that implements a plurality of stages each comprising one or more pipeline units, each pipeline unit configured to perform one or more operations on a block of pixels from a frame passing through the pipeline;

wherein one or more of the pipeline units in the block processing pipeline are each configured to:

process a current block at a unit core of the pipeline unit according to a current

configuration;

determine a configuration for processing a next block at the unit core according to
obtained information for the next block;

signal the unit core that the configuration for processing the next block is ready;

5

and

in response to said signal, process the next block at the unit core according to the
configuration for the next block;

wherein the pipeline unit determines the configuration for processing the next
block at the unit core while the unit core of the pipeline unit processes the
current block.

10

14. The device as recited in claim 13, wherein the information for the next block
includes information received from an upstream processing unit of the block processing pipeline
and information from a block that was previously processed by the processing unit.

15

15. The device as recited in claim 13 or 14, wherein each of the one or more pipeline
units is further configured to store the configuration for processing the next block at the unit core
to a configuration memory, wherein the unit core is configured to:

read the configuration from the configuration memory in response to said signal; and

20

signal to the pipeline unit that the configuration memory is clear.

1 / 19




FIG. 1
Prior Art

2 / 19




FIG. 2
Prior Art

3 / 19




FIG. 3

4 / 19

block processing pipeline 300




FIG. 4A




FIG. 4B

5 / 19




FIG. 4C

6 / 19




FIG. 5

7 / 19




FIG. 6

8 / 19




FIG. 7




FIG. 8A




FIG. 8B

10 / 19






FIG. 9



12 / 19




FIG. 11A




FIG. 11B

13 / 19




FIG. 12




FIG. 13

14 / 19




FIG. 14

15 / 19




FIG. 15A




FIG. 15B




FIG. 15C

16 / 19




FIG. 16

17 / 19




FIG. 17

18 / 19




FIG. 18

19 / 19




FIG. 19

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2014/051799

A. CLASSIFICATION OF SUBJECT MATTER
INV. G06T1/20
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 G06K H04N

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

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

EPO-Internal, INSPEC, 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 2010/135383 A1 (ZHAO WEIDONG [US] ET AL) 3 June 2010 (2010-06-03)	1,5
A	abstract figures 1,4,7 paragraphs [0001], [0005] - [0009], [0022], [0049] - [0053], [0056]	2-4,6-9
X	EP 2 161 685 A2 (SONY CORP [JP]; SONY ELECTRONICS INC [US]) 10 March 2010 (2010-03-10) abstract figures 1,3,7 paragraphs [0001], [0030], [0070] - [0076], [0104]	1,5



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents :

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

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

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

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

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

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

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

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

"&" document member of the same patent family

Date of the actual completion of the international search

17 November 2014

Date of mailing of the international search report

20/01/2015

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

Scholz, Volker

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2014/051799

Box No. II Observations where certain claims were found unsearchable (Continuation of item 2 of first sheet)

This international search report has not been established in respect of certain claims under Article 17(2)(a) for the following reasons:

1. ☐ Claims Nos.:
because they relate to subject matter not required to be searched by this Authority, namely:
2. ☐ Claims Nos.:
because they relate to parts of the international application that do not comply with the prescribed requirements to such an extent that no meaningful international search can be carried out, specifically:
3. ☐ Claims Nos.:
because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a).

Box No. III Observations where unity of invention is lacking (Continuation of item 3 of first sheet)

This International Searching Authority found multiple inventions in this international application, as follows:

see additional sheet

1. ☐ As all required additional search fees were timely paid by the applicant, this international search report covers all searchable claims.
2. ☐ As all searchable claims could be searched without effort justifying an additional fees, this Authority did not invite payment of additional fees.
3. ☐ As only some of the required additional search fees were timely paid by the applicant, this international search report covers only those claims for which fees were paid, specifically claims Nos.:
4. ☒ No required additional search fees were timely paid by the applicant. Consequently, this international search report is restricted to the invention first mentioned in the claims; it is covered by claims Nos.:

1-9

Remark on Protest

- ☐ The additional search fees were accompanied by the applicant's protest and, where applicable, the payment of a protest fee.
- ☐ The additional search fees were accompanied by the applicant's protest but the applicable protest fee was not paid within the time limit specified in the invitation.
- ☐ No protest accompanied the payment of additional search fees.

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2014/051799

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2010135383	A1	03-06-2010	NONE

EP 2161685	A2	10-03-2010	AU 2009213013 A1 25-03-2010
			CA 2678240 A1 09-03-2010
			CN 101673391 A 17-03-2010
			EP 2161685 A2 10-03-2010
			JP 2010067276 A 25-03-2010
			US 2010060651 A1 11-03-2010

FURTHER INFORMATION CONTINUED FROM PCT/ISA/ 210

This International Searching Authority found multiple (groups of) inventions in this international application, as follows:

1. claims: 1-9

a block processing pipeline comprising a software and a hardware pipeline

2. claims: 10-12

a block processing pipeline with signaling that the configuration for a block is ready

3. claims: 13-15

a block processing pipeline comprising pipeline units
