



(51) International Patent Classification:
H04N 7/26 (2006.01)

(21) International Application Number:
PCT/US2010/033744

(22) International Filing Date:
5 May 2010 (05.05.2010)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
12/436,196 6 May 2009 (06.05.2009) US

(71) Applicant (for all designated States except US): **ADVANCED MICRO DEVICES, INC.** [US/US]; One AMD Place, P.O. Box 3453, Sunnyvale, California 94088 (US).

(72) Inventor; and

(75) Inventor/Applicant (for US only): **FRANK, Michael** [DE/US]; 1325 Yarmouth Terrace, Sunnyvale, California 94087 (US).

(74) Agent: **MEYERTONS, HOOD, KIVLIN, KOWERT & GOETZEL, P.C.**; KIVLIN, B. Noel, 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, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, RO, RS, RU, 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, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, 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, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, 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))

(54) Title: EXECUTION UNITS FOR CONTEXT ADAPTIVE BINARY ARITHMETIC CODING (CABAC)




FIG. 1

(57) Abstract: A digital processor for recovering a source bitstream from an encoded bitstream that has been encoded according to a context adaptive binary arithmetic coding (CABAC) algorithm. The processor includes a first execution unit and a second execution unit. The first execution unit generates first execution data by operating on a first register and a second register, and stores the first execution data in the first register. The first execution data includes a current output bit, a temporary range value and a temporary offset value. The current output bit corresponds to a bit of the source bitstream. The second execution unit generates second execution data by operating on the first register and the second register, and stores the second execution data in the second register. The second execution data includes a normalized range value and a normalized offset value.



ARITHMETIC CODING (CABAC)

FIELD OF THE INVENTION

5 [0001] The present invention relates to the field of data compression, and more particularly, to mechanisms for performing encoding and/or decoding according to the context-adaptive binary arithmetic coding (CABAC) algorithm.

DESCRIPTION OF THE RELATED ART

10 [0002] Digital video encoding and decoding has become an important part of the workload for modern processors. Software based video decoding for MPEG2 at standard TV resolution (SD) has been around for almost 10 years, and HW acceleration for highly parallelizable parts of the decoding algorithm has been implemented in diverse graphic accelerators. With the introduction of more sophisticated video compression algorithms (MPEG4 and MPEG4/AVC, also known as
15 H.264) and even more so with the transition from SD to high-definition (HD) video streams, performance requirements for SW decoders have increased dramatically. (AVC is an acronym for “Advanced Video Coding”.) High definition video encoding (H.264) and JPEG2000 (now used for digital cinema) use arithmetic coding to losslessly reduce the data rate of the token stream generated by the lossy compression at the front end. Using arithmetic encoding is known
20 in the art to be optimal if the probability density function (PDF) of the encoded symbol stream is known. H.264 uses a variant of arithmetic coding, i.e., context-adaptive binary arithmetic coding (CABAC), that models the PDF (using a complex state machine) as a function of the current encoding context. Evaluating the incoming bitstream at high data rates and updating the context state is creating a huge processing load for current CPUs and consumes a significant part of the
25 performance of a processor core (e.g., greater than 50%). Thus, there exists a need for CPUs capable of performing the operations associated with CABAC encoding and/or decoding with increased efficiency.

SUMMARY

30 [0003] In one set of embodiments, a digital processor may be configured for decoding an input bitstream that has been encoded according to a context adaptive binary arithmetic coding (CABAC) algorithm. The processor may include a first execution unit, a second execution unit and a distribution unit.

[0004] The distribution unit may be coupled to the first execution unit and the second execution unit, and configured to fetch a stream of instructions, where the stream of instructions includes a first instruction and a second instruction. The distribution unit may be configured to: decode the first instruction and the second instruction; schedule execution of the decoded first instruction on
5 the first decode unit; and schedule execution of the decoded second instruction on the second decode unit.

[0005] In response to the distribution unit's scheduling of execution of the decoded first instruction, the first execution unit may be configured to generate a current output bit, a temporary range value and a temporary offset value based at least on a current range value and a
10 current offset value. The current offset value includes one or more bits of the input bitstream.

[0006] In response to the distribution unit's scheduling of execution of the decoded second instruction, the second execution unit may be configured to generate an update for the current range value and an update for the current offset value by normalizing respectively the temporary range value and the temporary offset value.

[0007] In some embodiments, the first execution unit is configured to operate on data fields in a
15 first register and data fields in a second register. The data fields in the first register include a binary most-probable-symbol (MPS) value, two next probability state values, and a set of least-probable-symbol (LPS) range values. The data fields in the second register include the current range value and the current offset value. The first register may include additional data fields that
20 are used to pass intermediate results from the first execution unit to the second execution unit, where the intermediate results include the temporary range value and the temporary offset value.

[0008] In addition, the first execution unit may be configured to generate a count of a number of leading zeros in the temporary range value. The leading zero count may be included as one of the intermediate results. The second execution unit may include a first bit shifting unit configured to
25 normalize the temporary range value by shifting the temporary range value based on the leading zero count.

[0009] In some embodiments, the first execution unit may include a first multiplexer, a second multiplexer and a subtraction unit. The first multiplexer may be configured to select one of the LPS range values as a temporary LPS range value, based on a selection index that is determined
30 by the current range value. The second multiplexer may be configured to generate a data value by selecting one of the constant zero, the constant two or the temporary LPS range value based on control information provided in the first register, wherein the control information includes a bypass mode control bit and a terminate mode control bit. The subtraction unit may be

configured to compute an intermediate range value from the current range value and the data value.

[0010] In some embodiments, the first execution unit may be configured to access memory to obtain updated context state information and to store the updated context state information in the second register.

[0011] In one set of embodiments, a method for performing CABAC decoding may involve the following actions. First, the method may involve executing a first instruction on a first execution unit of a microprocessor. The action of executing the first instruction may include generating a current output bit, a temporary range value and a temporary offset value based at least on a current range value and a current offset value, where the current offset value includes one or more bits of the input bitstream. The method may also involve executing a second instruction on a second execution unit of the microprocessor. The action of executing the second instruction may include generating an update for the current range value and an update for the current offset value by normalizing respectively the temporary range value and the temporary offset value.

BRIEF DESCRIPTION OF THE DRAWINGS

A better understanding of the present invention can be obtained when the following detailed description of the preferred embodiments is considered in conjunction with the following drawings.

Figure 1 illustrates one embodiment of a processor according to the present invention.

Figure 2A illustrates one embodiment of the data flow in and out of the execution unit 120.

Figure 2B illustrates one embodiment of the data flow in and out of the execution unit 130.

Figure 3 illustrates one embodiment of the data items contained in registers REG1 and REG2.

Figures 4A and 4B illustrates one embodiment of execution unit 120.

Figure 5 illustrates one embodiment of execution unit 130.

Figure 6 illustrates one embodiment of circuitry for computing the current output value (binVal), the temporary range value (tmpRange) and the temporary offset value (tmpOffset).

Figure 7 illustrates one embodiment of processor 100, includes execution units EU-1 through EU-N.

Figure 8 illustrates another embodiment of processor 100, i.e., a refinement of the embodiment shown in Figure 7.

Figure 9 illustrates an example instruction stream that may be supplied to processor 100 at least in some embodiments.

Figure 10 illustrates one embodiment of a processor 1000 that uses two registers (REG1 and REG2).

5 Figure 11 is a flowchart illustrating the CABAC decoding process.

Figure 12 is a flowchart showing a bypass decoding process.

Figure 13 is a flowchart showing a termination decoding process.

Figure 14 is a flowchart showing an arithmetic decoding process.

10 Figure 15 is a composite flowchart showing the CABAC decoding process as a whole according to one embodiment.

Figure 16 illustrates one embodiment of the context state.

Figure 17 illustrates one embodiment of the shared state.

Figure 18 illustrates one embodiment of the decoder state.

15 Figure 19 is pseudo code illustrating the operations performed by execution unit 120 according to one embodiment.

Figure 20 is pseudo code illustrating the operations performed by execution unit 130 according to one embodiment.

Figure 21 illustrates one embodiment of a method for performing CABAC decoding.

20 While the invention is susceptible to various modifications and alternative forms, 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.

25

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

30 [0012] Context adaptive binary arithmetic coding (CABAC) is a form of entropy coding that is used in the H.264 video compression standard. CABAC relies on context-adaptive models of the processed bitstream in combination with simple arithmetic operations. CABAC, as specified by the H.264 standard, requires a collection of 399 contexts. Each of the contexts stores information regarding the state of a corresponding finite state machine (FSM). It should be noted, however, that the principles of the present invention naturally generalize to any number of contexts.

[0013] In one set of embodiments, a processor (e.g., a CPU) may be configured to support software decoding and encoding according to the CABAC algorithm. Specifically, the processor may support specialized instructions that understand, interpret and concurrently update the complex state of the CABAC state machine. The lossless entropy encoding/decoding step in the H.264 video standard has been identified as being a significant bottleneck in the CABAC algorithm. The CABAC algorithm achieves very good compression but is essentially a serial algorithm that produces one bit of the data stream at a time.

[0014] In some embodiments, a digital processor 100 (e.g., a microprocessor) may be configured as shown in Figure 1 for decoding an input bitstream. The input bitstream is assumed to have been encoded according to a context adaptive binary arithmetic coding (CABAC) algorithm. The processor 100 may include a distribution unit 110, an execution unit 120 and an execution unit 130.

[0015] The distribution unit 110 may be coupled to the execution units 120 and 130, and configured to fetch a stream of instructions from memory. The stream of instructions includes instructions J_1 and J_2 , to be described below. The distribution unit may be further configured to: decode the instructions J_1 and J_2 to respectively generate decode results R_1 and R_2 ; schedule execution of the decode results R_1 on the execution unit 120; and schedule execution of the decode results R_2 on the execution unit 130.

[0016] In response to the distribution unit's action of scheduling execution of the decode result R_1 , the execution unit 120 may be configured to generate a current output bit (binVal), a temporary range value (tmpRange) and a temporary offset value (tmpOffset), based at least on a current range value (Range) and a current offset value (Offset). The current offset value may include one or more bits of the input bitstream. The current output bit represents the current result of the CABAC decoding algorithm.

[0017] In response to the distribution unit's action of scheduling execution of the decode result R_2 , the execution unit 130 may be configured to generate an update for the current range value (Range) and an update for the current offset value (Offset) by normalizing the temporary range value (tmpRange) and the temporary offset value (tmpOffset), respectively.

[0018] In some embodiments, the execution unit 120 may be configured to operate on data fields in a first register REG1 and data fields in a second register REG2, e.g., as shown in Figure 2A. See also Figure 3 for an illustration of one embodiment of registers REG1 and REG2. The data fields in the first register REG1 may include a binary most-probable-symbol (MPS) value, two next probability state values (nextStateMPS and nextStateLPS), and a set of least-probable-

symbol range values (e.g., an array rLPS[3:0] containing four values). The data fields in the second register REG2 may include the current range value (Range) and the current offset value (Offset).

[0019] The first register REG1 may also include additional data fields that are used to pass intermediate results from the execution unit 120 to the execution unit 130. (The intermediate results are also referred to herein as “shared state”.) The intermediate results may include the temporary range value (tmpRange) and the temporary offset value (tmpOffset). In Figure 3, the data field corresponding to the temporary range value (tmpRange) is denoted “NextRange” to suggest that it will be used in the next processing stage, i.e., by execution unit 130. Similarly, the data field corresponding to the temporary offset value (tmpOffset) is denoted “NextOffset”.

[0020] The execution unit 130 may be configured to operate on the intermediate results in the first register REG1 and on data fields in the second register REG2, and to generate results that are stored into the second register REG2, e.g., as illustrated in Figure 2B.

[0021] The processor 100 may include an architected set of registers specialized for SIMD instructions. (SIMD is an acronym for “simultaneous instruction multiple data”.) The registers REG1 and REG2 may be any two of the registers from the architected set, e.g., two registers specified by a current program instruction. In one embodiment, the architected set of registers is the set of 128-bit xmm registers as defined by the x86 instruction set.

[0022] Figures 4A and 4B illustrate one embodiment of execution unit 120. Figure 5 illustrates one embodiment of execution unit 130.

[0023] The execution unit 120 may be configured to generate a count of the number of leading zeros in the temporary range value, e.g., as illustrated by units 432, 440, 444, 450 and 452 of Figure 4B. The leading zero count (denoted by the variable SCount) may be stored in the first register REG1 as one of the intermediate results.

[0024] The execution unit 130 may include a first bit shifting unit configured to normalize the temporary range value (tmpRange) by shifting the temporary range value based on the leading zero count. For example, Figure 5 shows execution unit 130 as including a shift unit 516 to serve that normalization function.

[0025] In some embodiments, the execution unit 120 may include a multiplexer 414, a multiplexer 416 and a subtraction unit 418 as shown in Figure 4A. The multiplexer 414 may be configured to select one of the LPS range values as a temporary LPS range value (RangeLPS) based on a selection index (RangeIDX) that is determined from the current range value (Range). For example, in one embodiment, the selection index is the two most significant bits of the

current range value. The multiplexer 416 may be configured to generate a data value x by selecting one of the constant zero, the constant two or the temporary LPS range value based on control information provided in the first register REG1. The control information may include a bypass mode control bit (ByPass) and a terminate mode control bit (Terminate). The subtraction unit 418 may be configured to compute an intermediate range value (codIRange) from the current range value (Range) and the data value x.

[0026] In some embodiments, the execution unit 120 includes a subtraction unit 420, a logic unit LU₁ and a multiplexer 436 as shown in Figure 6. (Figures 4A and 4B illustrate one such embodiment, where the logic unit LU₁ is realized by units 426, 430 and 438.) The subtraction unit 420 may be configured to compute the temporary offset value (tmpOffset) by subtracting the intermediate range value (codIRange) from the current offset value (Offset). The intermediate range value equals an MPS range value (RangeMPS) in a normal decoding mode, equals the current range value (Range) in a bypass mode, and equals the current range value minus two in a terminate mode. The logic unit LU₁ may be configured to generate the current output bit (binVal) based on the MPS value, on control information provided in the first register REG1, and on a carry output bit (CY) of the subtraction unit 420. The current output bit represents a current result of the CABAC decoding algorithm. The multiplexer 436 may be configured to generate the temporary range value (tmpRange) by selecting either the intermediate range value (codIRange) or a temporary LPS range value (RangeLPS) based at least on the carry output bit (CY) of the subtraction unit 420. In some embodiments, additional logic intervenes between the carry output CY of the subtraction unit 420 and the select input of the multiplexer 436.

[0027] In some embodiments, the execution unit 120 includes a subtraction unit 420 and a multiplexer 422 as shown in Figures 4A and 4B. The subtraction unit 420 may be configured to compute the temporary offset value (tmpOffset) from the current offset value (Offset) and an intermediate range value (codIRange). The intermediate range value (codIRange) equals an MPS range value (RangeMPS) in a normal decoding mode, equals the current range value (Range) in a bypass mode, and equals the current range value minus two in a terminate mode. The multiplexer 422 may be configured to generate a current probability state value (tmpState) by selecting one of the two next probability state values (nextStateMPS or nextStateLPS) based on a carry output bit (CY) of the subtraction unit 420. The execution unit 120 may be configured to generate updates for the two next probability state values and the set of LPS range values (rLPS[3:0]) by accessing memory using an address that is based on the current probability state value (tmpState). The current probability state value may be stored in the CurrentState field of the register REG1.

[0028] In some embodiments, the second register REG2 may also include a buffer count value (BitsLeft), a first stream buffer (SBuf1) and a second stream buffer (SBuf2) as shown in Figures 3 and 5. The first and second stream buffers store bits of the input bitstream. The buffer count value (BitsLeft) indicates the number of valid bits remaining in the first stream buffer (SBuf1).

5 The execution unit 130 may include a bit shift unit 518 and a bit shift unit 522 as shown in Figure 5. The bit shift unit 518 may be configured to shift a data word DW based on the leading zero count (SCount) to generate a shifted data word (tmpVal). The data word DW may be a concatenation of the contents of the first stream buffer (SBuf1) and the temporary offset value (tmpOffset). The bit shift unit 522 may be configured to shift contents of the second stream
10 buffer (SBuf2) based on a shift control value (LCount) to generate a temporary buffer value (tmpBuf).

[0029] In some embodiments, the execution unit 130 includes a multiplexer 526 and an adder 530 as shown in Figure 5. The multiplexer 526 may be configured to generate a second temporary buffer value (tmpBuf2) by selecting the temporary buffer value (tmpBuf) or the
15 constant zero based on a control signal (tmpGet) that indicates whether the buffer count value (BitsLeft) is less than or equal to the leading zero count (SCount). The adder 530 may be configured to generate an output data word by adding the shifted data word (tmpVal) and the second temporary buffer value (tmpBuf2). The output data word includes a first portion that specifies the update for the current offset value (Offset) and a second portion that specifies an
20 update for the first stream buffer (SBuf1).

[0030] In some embodiments, the instruction J_1 (referred to above) may include information identifying the registers REG1 and REG2 from a set of architected registers in the digital processor. For example, the identifying information may identify the first register REG1 as both a source and destination operand for the instruction J_1 , and identify the second register REG2 as
25 an additional source operand for the instruction J_1 . The instruction J_1 may also include a memory address that points to a table stored in memory. The table will be described in more detailed below.

[0031] Furthermore, the instruction J_2 (referred to above) may include information identifying the registers REG1 and REG2 from the set of architected registers. For example, the identifying
30 information may identify the second register REG2 as a source and destination operand for the second instruction J_2 , and identify the first register REG1 as an additional source operand for the second instruction J_2 .

[0032] In some embodiments, the instruction stream received by the distribution unit 110 also includes instructions from a baseline instruction set B in addition to the instructions J₁ and J₂. Accordingly, the processor 100 may also include a set of execution units EU-1, EU-2, ..., EU-N configured for executing the instructions of the baseline instruction set, as shown in Figure 7.

5 The distribution unit 110 may be configured to decode instructions of the baseline instruction set and schedule execution of the decoded baseline instructions on the execution units EU-1 through EU-N.

[0033] As used herein, the term "baseline instruction set" is any instruction set that includes at least a set of general-purpose processing instructions such as instructions for performing integer and floating-point arithmetic, logic operations, bit manipulation, branching and memory access.

10 A "baseline instruction set" may also include other instructions, e.g., instructions for performing simultaneous-instruction multiple-data (SIMD) operations on integer vectors and/or on floating point vectors.

[0034] In some embodiments, the baseline instruction set may include an x86 instruction set such as the IA-32 instruction set from Intel or the AMD-64™ instruction set defined by AMD. In other embodiments, the baseline instruction set may include the instruction set of a processor such as a MIPS processor, a SPARC processor, an ARM processor, a PowerPC processor, etc. The baseline instruction set may be defined in an instruction set architecture.

15

[0035] Figure 8 illustrates one embodiment of the processor 100. In addition to the units shown in Figure 7, this embodiment includes an instruction cache 105, a load/store unit 140, a register file 160 and a data cache 170.

20

[0036] The instruction cache 105 may store copies of instructions that have been recently accessed from system memory. (System memory resides external to processor 100.) The distribution unit 110 may fetch the stream S of instructions from the instruction cache 105. The stream S includes instructions drawn from the baseline instruction set (as described above) and from an extension set. The extension set includes the instructions J₁ and J₂, as described above.

25

[0037] The programmer has the freedom to intermix instructions of the baseline instruction set and the extension set when building a program for processor 100. An example of this mixing of instructions is shown in Figure 9. Example stream 300 includes instructions I₀, I₁, I₂, I₃, I₄, I₅, ... from the baseline instruction set, and instructions J₁ and J₂ from the extension set. While the example stream 300 shows one occurrence of the pair J₁ and J₂, that instruction pair may occur any number of times. For example, a long video stream may involve many occurrences of the pair J₁ and J₂.

30

[0038] The distribution unit 110 may decode the stream of fetched instructions into executable operations (ops). Each fetched instruction may be decoded into one or more ops. Some of the fetched instructions (e.g., some of the more complex instructions) may be decoded by accessing a microcode ROM. Furthermore, some of the fetched instructions may be decoded in a one-to-one fashion, i.e., so that the instruction results in a single op that is unique to that instruction.

[0039] The distribution unit 110 may schedule the ops for execution on the execution units that include: execution units 120 and 130, execution units EU-1 through EU-N, and load/store unit 140. The one or more ops (R_1) that result from decoding the instruction J_1 are scheduled for execution on execution unit 120, and the one or more ops (R_2) that result from decoding the instruction J_2 are scheduled for execution on execution unit 130.

[0040] As noted above, the distribution unit 110 may decode each instruction of the stream S of fetched instructions into one or more ops and schedule the one or more ops for execution on appropriate ones of the executions units. In some embodiments, the distribution unit 110 may be configured for superscalar operation, out-of-order (OOO) execution, multi-threaded execution, speculative execution, branch prediction, or any combination thereof. Thus, in various embodiments, distribution unit 110 may include various combinations of: logic for determining the availability of the execution units; logic for dispatching two or more ops in parallel (in a given clock cycle) whenever two or more execution units capable of handling those ops are available; logic for scheduling the out-of-order execution of ops and guaranteeing the in-order retirement of ops; logic for performing context switching between multiple threads and/or multiple-processes; etc.

[0041] Load/store unit 140 may couple to a data cache 170 and may be configured to perform memory write and memory read operations. For a memory write operation, the load/store unit 140 may generate a physical address and the associated write data. The physical address and write data may be entered into a store queue (not shown) for later transmission to the data cache 170. Data read from memory may be supplied to load/store unit 140 from data cache 170 (or from an entry in the store queue in the case of a recent store).

[0042] Execution units EU-1 through EU-N may include one or more integer pipelines and one or more floating-point units. The one or more integer pipelines may include resources for performing integer operations (such as add, subtract, negate, multiply and divide), logic operations (such as AND, OR and XOR), and bit manipulation (such as shift and cyclic shift). In some embodiments, resources of the one or more integer pipelines are configured to perform SIMD integer operations. The one or more floating-point units may include resources for

performing floating-point operations. In some embodiments, the resources of the one or more floating-point units are configured to perform SIMD floating-point operations.

[0043] In one set of embodiments, the execution units EU-1 through EU-N may include one or more SIMD units configured for performing integer and/or floating point SIMD operations.

5 [0044] As illustrated by Figure 8, the execution units may couple to a dispatch bus 122 and a results bus 150. The execution units receive ops from the distribution unit 110 via the dispatch bus 122, and pass the results of execution to register file 160 via results bus 150. The register file 160 couples to feedback path 158, which allows data from the register file 160 to be supplied as source operands to the execution units. Bypass path 155 couples between results bus 150 and the
10 feedback path, allowing the results of execution to bypass the register file 160, and thus, to be supplied as source operands to the execution units more directly. Register file 160 may include physical storage for a set of architected registers.

[0045] As noted above, the execution units EU-1 through EU-N may include one or more floating-point units. Each floating-point unit may be configured to execute floating-point
15 instructions (e.g., x87 floating-point instructions, or floating-point instructions compliant with IEEE 754/854). Each floating-point unit may include an adder unit, a multiplier unit, a divide/square-root unit, etc. Each floating-point unit may operate in a coprocessor-like fashion, in which case the distribution unit 110 directly dispatches the floating-point instructions to the floating-point unit. The floating-point unit may include storage for a set of floating-point
20 registers (not shown).

[0046] In some embodiments, the processor 100 may be included in a computer system such as a personal computer (in any of its various forms), a server computer, a client computer, a mobile device, a digital television, a set-top box, etc. The processor 100 may be implemented in terms of one or more integrated circuits.

25 [0047] In one set of embodiments, a digital processor 1000 for recovering a source bitstream from a CABAC encoded bitstream may be configured as shown in Figure 10. The processor 1000 may include an execution unit 1020 and an execution unit 1030. One embodiment of execution unit 1020 is shown in Figures 4A and 4B, and one embodiment of execution unit 1030 is shown in Figure 5.

30 [0048] The execution unit 1020 may be configured to generate execution data by operating on a first register REG1 and a second register REG2, and to store the execution data in the first register REG1. (The two occurrences of the first register REG1 shown in Figure 10 are for diagrammatic simplicity, not to imply that the first register is duplicated in hardware. Similarly,

the three occurrences of the second register REG2 are for diagrammatic simplicity.) The execution data may include a current output bit (binVal), a temporary range value (tmpRange) and a temporary offset value (tmpOffset), where the current output bit corresponds to a bit of the source bitstream.

5 [0049] The execution unit 1030 may be configured to generate second execution data by operating on the first register REG1 and the second register REG2 and to store the second execution data in the second register REG2. The second execution data includes a normalized range value (Range) and a normalized offset value (Offset).

10 [0050] The first register REG1 may include data fields for storing the current output bit (binVal), the temporary range value (tmpRange), the temporary offset value (tmpOffset), a binary most-probable-symbol (MPS) value, a current probability state (CurrentState), two next probability state values (nextStateMPS and nextStateLPS), and a set of least-probable-symbol (LPS) range values, e.g., as shown in Figure 3. The first register REG1 may also include data fields for storing a bypass control bit, a termination control bit and a zero state indicator. The bypass
15 control bit indicates whether a bypass mode is selected. The termination control bit indicates whether a terminate mode is selected. The zero state indicator indicates whether the probability state index of the current context is equal to zero.

[0051] In some embodiments, the execution data generated by the execution unit 1020 also includes an update for the current probability state (CurrentState) and an update for the MPS
20 value. Furthermore, the execution unit 1020 may be configured to initiate a memory access and to store results of the memory access in the first register REG1. The results of the memory access may include updates for the two next probability state values and for the set of LPS range values.

[0052] The processor 1000 may include additional execution units configured to execute instructions of a baseline instruction set, e.g., as described above and/or as shown in Figure 8. In
25 one embodiment, the baseline instruction set includes a set of integer arithmetic operations and a set of floating point arithmetic operations.

[0053] In some embodiments, the second register REG2 includes data fields for storing a first stream buffer (SBuf1), the normalized range value (Range) and the normalized offset value (Offset), e.g., as shown in Figure 3. The first stream buffer (SBuf1) stores bits of the encoded bit
30 stream. Furthermore, the execution unit 1030 may be configured to extract a number of bits from the first stream buffer (SBuf1), and to generate the normalized offset value by left shifting the temporary offset value (tmpVal) and injecting the extracted bits into the LSB positions vacated by the left shifting.

[0054] The second register REG2 may also include a remaining bit count (BitsLeft) that indicates a number of valid bits remaining in the first stream buffer (SBuf1).

[0055] In one set of embodiments, a processor such as processor 100 or processor 1000 may be configured based on a reorganization of the complex state of the CABAC decoding process so that part of the state is stored within two registers (e.g., wide registers such as the 128-bit wide xmm registers). This state may be manipulated using the instructions J1 and J2 as described herein. The instructions J₁ and J₂ may interact with each other by communicating intermediate state via predefined locations within a register so that the repeated execution of the instruction pair {J₁, J₂} decodes at least one bit of data per iteration. The processor may in some embodiments be configured according to a global storage organization for the context state that enables the processor to load the entire context state into one of the registers in a single memory access and store the updated context state into memory in a single memory access whenever a context change is required.

[0056] In some embodiments, the processor may be configured to support multi-tasking. Thus, two or more streams of CABAC data may be decoded in parallel (or apparently in parallel from the point of view of the user/users). The mechanisms for saving state on task switches are well understood and supported by any of various operating systems.

[0057] In some embodiments, the processor includes two or more independent cores, where each core includes logic for executing the instructions J1 and J2. Thus, two or more streams of CABAC data may be decoded in parallel.

[0058] In some embodiments, the processor may be configured so that the instructions J₁ and J₂ execute in parallel (or, at least partially in parallel). For example, execution unit 120 may still be accessing memory for the information to update the context state while the second execution unit 130 is operating on the decoder state and the shared state. (See Figures 4A, 4B and 5.)

[0059] CABAC decoding is defined in section 9.3.3.2 ("Arithmetic Decoding Process") of the ITU-T Recommendation H.264, entitled "Advanced Video Coding for Generic Audiovisual Services", published in May 2003 by the International Telecommunication Union (ITU). That document, hereafter referred to as "the H.264 Specification", is incorporated by reference in its entirety. Section 9.3.3.2 and subsequent sections describe in detail the decoding process for a binary symbol. See Figure 11 for a flowchart of the decoding process 1100.

[0060] Arithmetic coding is based on the principle of recursive subdivision of an initial interval of the real line. Suppose the binary symbols 0 and 1 have respective probabilities of occurrence equal to p(0) and p(1)=1-p(0). An initial code interval having range codIRange is divided into

two subintervals of range $p(0)*codIRange$ and $(1-p(0))*codIRange$, respectively. (The term “range” is used in the sense of “length”.) Depending on the current input value, the decoding process 1100 returns the binary decision 0 or 1, and updates the code interval to determine a new code interval. The binary decision is the output of the decoding process and hopefully is equal to the corresponding bit of the original bit stream that was encoded at the corresponding transmitter/encoder.

[0061] The decoding process 1100 distinguishes between the most probable symbol (MPS) and the least probable symbol (LPS). The MPS and LPS are complementary: $MPS=1-LPS$. However, the MPS (and thus, the LPS) changes over time since the probabilities $p(0)$ and $p(1)$ are not the same for every symbol (i.e., bit) of the original bit stream. If $p(0)$ is greater than $\frac{1}{2}$, then $MPS=0$. Otherwise, $MPS=1$. The binary decision of the decoding process 1100 depends on the current context, which is identified by the index $ctxIdx$.

[0062] In one embodiment, each context has a set of state data that includes: a bit specifying which symbol (0 or 1) is the MPS, and a 6-bit value representing the occurrence probability of the LPS. Thus, in this embodiment, each context has 128 states, i.e., the product of two states for the MPS and 64 states for the LPS probability. However, it is noted that the specific choice of the MPS and the LPS probability value as the context state data is not essential. More generally, the context state data may include either one of MPS or LPS, and, either one of MPS probability or LPS probability, implying at least four different embodiments for the definition of the context state data. Furthermore, the selection of 64 as the number of probability states is not essential. The principles of the present invention naturally generalize to any number of probability states.

[0063] The decoding process 1100 may have the following properties. First, the decoding process may maintain a finite state machine for each of the contexts. For a given context, the finite state machine may control the transitions between states of the corresponding context. As described above, the context state information may include one bit for MPS state and 6 bits for LPS probability. The 6-bit probability state index may map to probability values in a non-linear fashion, e.g., as defined by the following expression: $p_k = \alpha * p_{k-1}$ for $k = 1, 2, 3, \dots, 63$, with $\alpha = (0.01875/0.5)^{1/63}$ and $p_0 = 0.5$. The index values may map to probability values in a decreasing fashion, with index 0 corresponding to $pLPS=0.5$, and index 63 corresponding approximately to $pLPS=0.0$. It is noted that the principles of the present invention naturally generalize to a wide variety of other mappings between probability state index and probability value, including both linear and non-linear mappings.

[0064] Second, the range ($codIRange$) may be quantized to 2 bits (i.e., to one of 4 values).

[0065] Third, any syntax elements that are (or are expected to be) sufficiently approximated by a constant probability density function (PDF) may bypass the arithmetic decoding process, and thus, not impact the context FSMs.

[0066] The input to the decoding process may include a bypassFlag, an index (ctxIdx) into a context-state array, and the variables codIRange and codIOffset. The variables codIRange and codIOffset may be interpreted as the state of the decoding process. As shown in Figure 11, the decoding process 1100 includes a determination 1110 if the bypassFlag is TRUE. If so, the bypass decoding process 1116 is performed. Otherwise, control passes to operation 1112. In operation 1112, the context index ctxIdx is compared to a termination value. (In one embodiment, the termination value is 276. However, the exact value of the termination index is not essential. Any other value would suffice.) If the context index is equal to the termination value, a termination process 1120 is performed. Otherwise, an arithmetic decoding process 1114 is performed.

[0067] In one set of embodiments, a programmable processor may be designed based on an architectural plan that involves organizing the variables involved in the CABAC decoding process, compacting the data structures that hold the state, and parallelizing the computations used in the state transition process. In one embodiment, the state transition computations may be implemented in an MIMD (multiple-instruction multiple-data) fashion.

Bypass Decoding

[0068] As described above in connection with Figure 11, the arithmetic decoding process 1114 may be bypassed, and a bypass decoding process 1116 may be performed if the bypassFlag is equal to one. The bypass decoding process may take bits from the input bitstream and convert them into output data bits as shown in Figure 12.

[0069] At 1210, the internal variable codIOffset may be doubled, and a bit from the input bitstream may be injected (inserted) into the LSB position of that variable. The doubling may be achieved simply by shifting the variable contents one bit position to the left.

[0070] At 1212, a test is performed to determine if the variable codIOffset is greater than or equal to codIRange. If that variable is greater than or equal to codIRange, control passes to 1214. Otherwise, control passes to 1216.

[0071] At 1214, the binary decision binVal is set equal to one and the variable codIOffset is updated by subtracting codIRange from codIOffset: $\text{codIOffset} = \text{codIOffset} - \text{codIRange}$.

[0072] At 1216, the binary decision is set equal to zero.

[0073] The bypass decoding process 1116 is finished after 1214 or after 1216. However, in alternative embodiments, the bypass decoding process may include one or more additional operations.

Termination Decoding

[0074] As described above in connection with Figure 11, the termination decoding process 1120 may be performed if the bypass Flag is set to zero and the context index ctxIdx is equal to the termination value. The termination decoding process 1120 may be performed at the end of a slice. The termination decoding process 1120 may involve the operations shown in Figure 13.

[0075] At 1310, the variable codIRange is decremented by 2.

[0076] At 1312, a test is performed to determine if the variable codIOffset is greater than or equal to the variable codIRange. If the test indicates that codIOffset is greater than or equal to codIRange, control passes to 1314. Otherwise, control passes to 1316.

[0077] At 1314, the binary decision binVal is set equal to 1. After 1314, the termination decoding process 1120 is finished. However, other embodiments are contemplated where one or additional operations are included after 1314.

[0078] At 1316, the binary decision binVal is set equal to 0. At 1318, the variable codIRange is renormalized by shifting the contents of the variable codIRange to the left by the number of leading zeros. After 1318, the termination decoding process is finished. However, other embodiments are contemplated where one or additional operations are included after 1318.

Arithmetic Decoding Process for Binary Decision

[0079] The arithmetic decoding process 1114 may take the value codIOffset and codIRange as input values and return the decision value binVal as output. The arithmetic decoding process 1114 may determine the binary decision binVal by determining if CodIOffset is within the range of the least probable symbol (LPS) or within the range of the most probable symbol (MPS). As an optimization, the second input variable codIRange may be quantized to a two-bit value, and the exact value of the current LPS range (codIRangeLPS) may be looked up in a table that is accessed using the probability state index and the two-bit quantized range value. The table may represent the quantized PDF for each probability state. In one embodiment, the table is identical to Table 9-33 of the H.264 specification. (While the H.264 specification specifies a 2-bit quantization of the variable codIRange, alternative embodiments may use N-bit quantization with N different from two.)

[0080] As shown in Figure 14, the arithmetic decoding process 1114 includes a process of decoding the input data and a process of updating state information (i.e., both system state information and context state information).

[0081] At 1410, a quantized range value $qCodIRangeIdx$ is computed based on the variable $codIRange$, and then, the quantized range value is used to compute $codIRangeLPS$ (a range value for the LPS) and $codIRange$ (a tentative range value for the MPS). The quantized range value $qCodIRangeIdx$ is computed from the variable $codIRange$ based on the following relation: $qCodIRangeIdx = (codIRange \gg 6) \& 0x03$. The expression “ $A \gg B$ ” denotes a B bit shift to the right of the contents of operand A. The notation “ $\&$ ” denotes the logical AND operator. The notation “ $0x03$ ” denotes the number three represented as a 2-digit hexadecimal value.

[0082] The quantized range value $qCodIRangeIdx$ and the current probability state index $pStateIdx$ are used to look up the value $codIRangeLPS$ in the table $rangeTabLPS$. This table lookup is represented by the expression:

[0083] $codIRangeLPS = rangeTabLPS[pStateIdx][qCodIRangeIdx]$.

[0084] The variable $codIRange$ is updated by subtracting $codIRangeLPS$ from $codIRange$: $codIRange = codIRange - codIRangeLPS$.

[0085] At 1412, a test is performed to determine if the variable $codIOffset$ is greater than or equal to $codIRange$. If the test determines that $codIOffset$ is greater than or equal to $codIRange$, control passes to 1414. Otherwise, control passes to 1422.

[0086] At 1414, the binary decision $binVal$ is set equal to the complement of $valMPS$, i.e., $binVal = !valMPS$. (The variable $valMPS$ represents the value of the most probable symbol.) Furthermore, the variables $codIOffset$ and $codIRange$ are updated according to the relations: $codIOffset = codIOffset - codIRange$, and $codIRange = codIRangeLPS$.

[0087] At 1416, a test is performed to determine if the probability state index $pStateIdx$ is equal to zero. If $pStateIdx$ is determined to be equal to zero, control passes to 1418. Otherwise, control passes to 1420.

[0088] At 1418, the variable $valMPS$ is complemented: $valMPS = 1 - valMPS$. After 1420, control passes to 1420.

[0089] At 1420, an updated value for the probability state index $pStateIdx$ is determined by accessing a first state transition table ($transIdxLPS$) using the current value of $pStateIdx$. After 1420, control passes to 1424.

[0090] At 1422 (which is performed when test 1412 determines that $codIOffset$ is not greater than or equal to $codIRange$), the binary decision $binVal$ is set equal to $valMPS$, and an updated

value for the probability state index pStateIdx is determined by accessing a second state transition table (transIdxMPS) using the current value of pStateIdx. After 1422, control passes to 1424.

[0091] At 1424, the value codIRange may be normalized. The normalization may be performed according to the following pseudo code:

```

5         while (codIRange < 0x0100) {
                codIRange = codIRange << 1
                codIOffset = codIOffset << 1
                codIOffset = codIOffset | read_bits(1)
        }

```

10 [0092] The operator “|” represents bitwise OR. In other words, if codIRange is less than 0x0100, codIRange and codIOffset are left-shifted by one bit, and a bit from the input bitstream is injected in the LSB position of codIOffset. These operations may be repeated until the condition becomes true.

[0093] After 1424, the arithmetic decoding process 1114 is finished. In alternative embodiments, 15 one or one additional processes may be included in the arithmetic decoding process.

[0094] In one embodiment, the processor may maintain context state information for a current context (corresponding to ctxIdx). The context state information (ContextState) may include: a bit specifying the current MPS; a current probability state index (CurrentState); a next probability state index (NextStateMPS) under the assumption that the decoded symbols is equal to the MPS; 20 and a next probability state index (NextStateLPS) under the assumption that the decoded symbol is equal to the LPS. (The NextStateMPS and NextStateLPS fields may be populated with data based on Table 9-34 of the H.264 specification.) Thus, if the decoded symbol is equal to the MPS, the probability state transition is given by the expression:

[0095] ContextState.CurrentState=ContextState.NextStateMPS.

25 [0096] Otherwise, the probability state transition is given by the expression:

[0097] ContextState.CurrentState= ContextState.NextStateLPS.

[0098] Figure 15 shows a composite flowchart that results from substituting the flowcharts of Figure 12-14 into the high-level flowchart of Figure 11. The composite flowchart reveals that there is substantial commonality in the different branches of the data flow - considering that the 30 state information is shared between the branches. (The normalization operations in Figures 13 and 14 have been consolidated into normalization operation 1510.) As described above, the decoding process 1100 uses the flow control bit bypassFlag to control branching to the bypass decoding process (as shown at 1110). If bypassFlag is false, the context index (ctxIdx) is

compared to a termination value (as shown at 1112) to decide if the flow branches to the termination decoding process. Conveying all of this information to a single processor instruction as separate parameters may exceed the number of operands supported by many existing processor architectures. Thus, in some embodiments, the information may be consolidated into one or more registers, e.g., xmm registers.

[0099] In some embodiments, the context state may be extended to include space for storing an entry from rangeTabLPS (see Table 9-33 of the H.264 specification). Thus, the critical data used to determine the next subinterval may be prefetched.

[00100] The context state may be further extended with two entries representing the two probability states that are possible from the current probability state. (These two probability states are referred to as “next probability states”.) This allows the processor to perform the calculations required for a state transition after only a single access to the stored context state.

[00101] The context state may also include: a bit (TerminationSymbol) that is used to indicate if the given context has context index equal to the termination value, and another bit (bypassFlag) that signals selection of the bypass decoding process.

[00102] In one embodiment, bypassFlag is set only in the last context of the context array (context number 399 in embodiments conforming to the H.264 specification). In order to select the bypass decoding process, the context state of the last context is loaded into the processor.

[00103] In some embodiments, the processor may be configured to use a state register. The state register may be updated with a single access to the state transition table. This access may be at least partially hidden behind the second phase, which includes the renormalization operation, the data extraction and the determination of the next context state.

[00104] In some embodiments, a context table and a state transition table may be defined as specified in the following pseudo code.

```

struct {
    byte rLPS[3:0] ;
    state nextStateMPS;
    state nextStateLPS;
    state currentState;
    bit MPS;
    bit currentState_is_Zero;
    bit previousMPS;
    bit TerminationSymbol;

```

```

        bit bypassFlag;
    } ContextTable[0:399];

```

```

    struct {
        byte rLPS[3:0];
        state nextStateMPS;
        state nextStateLPS;
    } StateTable[0:63];

```

[00105] In one embodiment, the “state” data type is a 6-bit unsigned integer type and is used to define any of the probability state indices: currentState, nextStateMPS, nextStateLPS.

[00106] In some embodiments, state information may be stored in the 128-bit wide xmm registers. The xmm registers provide sufficient storage space for multiple heterogeneous data fields. This use of the xmm registers may solve the problem of updating multiple items in an instruction. The architectural definition for the registers may partition the state information into three types: context state, shared state and decoder state.

[00107] One embodiment of a structure named “ContextState”, for storing context state information, is shown in Figure 16. ContextState may be 64 bits wide and occupy the lower half of an xmm register. (See also register REG1 of Figure 3.) As shown in Figure 16, ContextState resembles a record of ContextTable so that a single 64-bit memory access can move an entire record from/to memory. Assuming a 512-bit cache line, eight of these context records can be packed into a cache line. Other cache line widths are contemplated as well.

[00108] One embodiment of a structure named “SharedState” is shown in Figure 17. SharedState may occupy the upper 64 bits of the same xmm register whose lower half is occupied by ContextState. SharedState holds intermediate results generated by the first phase of the decoder, i.e., by the execution unit 120. These intermediate results may be used by the second phase (i.e., by execution unit 130) or by software during the DecoderBypass operation.

[00109] One embodiment of a structure named DecoderState is shown in Figure 18. DecoderState may be 128 bits wide and occupy another xmm register. The upper 64-bits of the register may serve as a bitstream buffer. Data is automatically extracted from this buffer, and, on underrun of the lower 32 bits, a get_data flag is raised.

[00110] The mapping of data items to register fields shown in Figures 16-18 is meant as one possible example. A wide variety of other mappings are possible. In one alternative embodiment, the mapping of shared state and context state respectively to the upper and lower

halves of the same xmm register is reversed, i.e., shared state occupying the lower half and context state the upper half. Furthermore, the present invention imposes no requirement that context state and shared state be segregated. In one embodiment, they are interleaved (or partially interleaved).

5 [00111] The CABAC binary decode may be implemented using two instructions. The first instruction J_1 is responsible for updating the context state and generating intermediate state. The second instruction J_2 updates the decoder state, including the data buffer SBuf1.

[00112] The first instruction J_1 may involve three operands -- two source operands (src1 and src2) and one destination operand (dest). For example, according to one possible formatting
10 convention, the first instruction may have the following syntax:

CABACTX dest, src1, src2, mem64

[00113] The operands dest, src1 and src2 designate xmm registers with the constraint that dest=src1. The register src1 contains the context state and shared state information. The register src2 contains the decoder state information. The memory address mem64 is a pointer to the base
15 address of the state transition table that is stored in memory. The first instruction J_1 may be executed by execution unit 120, e.g., as shown in Figures 4A and 4B. According to one embodiment, the first instruction J_1 may operate as described by the pseudo code given in Figure 19.

[00114] The second instruction J_2 may also include three operands -- two source operands
20 (src1 and src2) and one destination operand (dest). For example, the second instruction may have the following syntax:

CABACSX dest, src2, src1

[00115] The operands dest, src1 and src2 designate xmm registers with the constraint that dest=src2. The register src2 holds the decoder state information. The lower 64 bits of register
25 src1 hold the context state information; the upper 64 bits provide shared state information. This second instruction J_2 may be executed by execution unit 130, e.g., as shown in Figure 5. According to one embodiment, the second instruction J_2 may operate as described by the pseudo code given in Figure 20.

[00116] In one alternative set of embodiments, the processor may be configured to
30 implement the instruction J_1 by splitting the 128-bit register that contains the context state and shared state into two 64-bit pieces and executing two ops (e.g., microinstructions) in a way that updates the lower half and the upper independently. Similarly, the second instruction J_2 may also

to be split into separate parts. The distribution unit 110 may be configured to implement the splitting of instructions J_1 and J_2 into ops.

[00117] In one set of embodiments, a method for performing CABAC decoding may involve the following actions, as shown in Figure 21. The method operates on an input bitstream
5 that has been encoded according to the CABAC algorithm.

[00118] At 2110, the method may involve executing a first instruction on a first execution unit of a microprocessor. The action of executing the first instruction may include generating a current output bit, a temporary range value and a temporary offset value based at least on a current range value and a current offset value, where the current offset value includes one or more
10 bits of the input bitstream.

[00119] At 2120, the method may also involve executing a second instruction on a second execution unit of the microprocessor. The action of executing the second instruction may include generating an update for the current range value and an update for the current offset value by normalizing respectively the temporary range value and the temporary offset value.

[00120] The input bitstream may be decoded by repeatedly performing 2110 and 2120, i.e., by repeatedly executing the first instruction and second instruction. One bit of the original source stream is generated each time the instruction pair is executed.

[00121] While much of the above description has focused on embodiments for decoding a bit stream according to the H.264 specification, the inventive principles described herein
20 naturally generalize other sets of data lengths, register lengths, parameter values, array sizes, etc.

[00122] The principles disclosed herein may be used to create a pair of execution units for encoding a bitstream according to the CABAC algorithm. The execution units may re-use the layout of the complex state (e.g., as shown in Figure 3) and implement the logic for the (now different) state transitions and the stream packing. This pair of execution units may be
25 configured to execute two new instructions K_1 and K_2 , respectively.

[00123] Although the embodiments above have been described in considerable detail, numerous variations and modifications will become apparent to those skilled in the art once the above disclosure is fully appreciated. It is intended that the following claims be interpreted to embrace all such variations and modifications.

1. A digital processor for decoding an input bitstream that has been encoded according to a context adaptive binary arithmetic coding (CABAC) algorithm, the processor comprising:

5 a first execution unit;

a second execution unit; and

a distribution unit coupled to the first execution unit and the second execution unit, and configured to fetch a stream of instructions, wherein the stream of instructions includes a first instruction and a second instruction, wherein the distribution unit is configured to:

10 decode the first instruction and the second instruction;

schedule execution of the decoded first instruction on the first decode unit; and

schedule execution of the decoded second instruction on the second decode unit;

wherein, in response to the distribution unit's scheduling of execution of the decoded first instruction, the first execution unit is configured to generate a current output bit, a temporary
15 range value and a temporary offset value based at least on a current range value and a current offset value, wherein the current offset value includes one or more bits of the input bitstream;

wherein, in response to the distribution unit's scheduling of execution of the decoded second instruction, the second execution unit is configured to generate an update for the current range value and an update for the current offset value by normalizing respectively the temporary range
20 value and the temporary offset value.

2. The digital processor of claim 1, wherein the first execution unit is configured to operate on data fields in a first register and data fields in a second register, wherein the data fields in the first register include a binary most-probable-symbol (MPS) value, two next probability state
25 values, and a set of least-probable-symbol (LPS) range values, wherein the data fields in the second register include the current range value and the current offset value, wherein the first register includes additional data fields that are used to pass intermediate results from the first execution unit to the second execution unit, wherein the intermediate results include the temporary range value and the temporary offset value.

3. The digital processor of claim 2, wherein the first execution unit is configured to generate a count of a number of leading zeros in the temporary range value, wherein the intermediate results also include the leading zero count, wherein the second execution unit includes a first bit

shifting unit configured to normalize the temporary range value by shifting the temporary range value based on the leading zero count.

4. The digital processor of claim 2, wherein the first execution unit includes:

a first multiplexer configured to select one of the LPS range values as a temporary LPS range value, based on a selection index that is determined by the current range value;
a second multiplexer configured to generate a data value by selecting one of the constant zero, the constant two or the temporary LPS range value based on control information provided in the first register, wherein the control information includes a bypass mode control bit and a terminate mode control bit; and
a first subtraction unit configured to compute an intermediate range value from the current range value and the data value.

5. The digital processor of claim 2, wherein the first execution unit includes:

a second subtraction unit configured to compute the temporary offset value from the current offset value and an intermediate range value, wherein the intermediate range value equals an MPS range value in a first decoding mode;
first logic configured to generate the current output bit based on the binary MPS value, on control information provided in the first register, and on a carry output bit of the second subtraction unit, wherein the current output bit represents a current result of the CABAC decoding algorithm;
a third multiplexer configured to generate the temporary range value by selecting either the intermediate range value or a temporary LPS range value based at least on the carry output bit of the second subtraction unit.

6. The digital processor of claim 2, wherein the first execution unit includes:

a second subtraction unit configured to compute the temporary offset value from the current offset value and an intermediate range value, wherein the intermediate range equals an MPS range value in a first decoding mode;
a fourth multiplexer configured to generate a current probability state value by selecting one of the two next probability state values based on a carry output bit of the second subtraction unit;

wherein the first execution unit is configured to generate updates for the two next probability state values and the set of LPS range values by accessing memory using an address based on the current probability state value.

5 7. The digital processor of claim 2, wherein the second register also includes a buffer count value, a first stream buffer and a second stream buffer, wherein the first and second stream buffers store bits of the input bitstream, wherein the buffer count value indicates the number of valid bits remaining in the first stream buffer, wherein the second execution unit includes:

10 a first bit shifting unit configured to shift a first data word based on the leading zero count to generate a first shifted data word, wherein the first data word is a concatenation of contents of the first stream buffer and the temporary offset value;

a second bit shifting unit configured to shift contexts of the second stream buffer based on a shift control value to generate a temporary buffer value.

15 8. The digital processor of claim 7, wherein the second execution unit includes:

a first multiplexer configured to generate a second temporary buffer value by selecting the temporary buffer value or the constant zero based on a control signal that indicates whether the buffer count value is less than or equal to the leading zero count;

20 an adder configured to generate an output data word by adding the first shifted data word and the second temporary buffer value, wherein the output data word includes a first portion that specifies the update for the current offset value and a second portion that specifies an update for the first stream buffer.

25 9. The digital processor of claim 2, wherein the first instruction includes first information identifying the first and second registers from a set of architected registers in the digital processor, wherein the first information identifies the first register as both a source and destination operand for the first instruction, wherein the first information also identifies the second register as an additional source operand for the first instruction, wherein the first instruction also includes a memory address that points to a table stored in memory.

30 10. The digital processor of claim 9, wherein the second instruction includes second information identifying the first and second registers from the set of architected registers, wherein the second information identifies the second register as a source and destination operand for the

second instruction, wherein the second instruction identifies the first register as an additional source operand for the second instruction.

11. The digital processor of claim 1 further comprising a set of two or more execution units, wherein the stream of instructions includes instructions of a baseline instruction set, wherein the distribution unit is configured to decode instructions of the baseline instruction set and schedule execution of the decoded instructions on the set of two or more executions units.

12. A digital processor for recovering a source bitstream from an encoded bitstream that has been encoded according to a context adaptive binary arithmetic coding (CABAC) algorithm, the processor comprising:

a first execution unit configured to generate first execution data by operating on a first register and a second register, and to store the first execution data in the first register, wherein the first execution data includes a current output bit, a temporary range value and a temporary offset value, wherein the current output bit corresponds to a bit of the source bitstream;

a second execution unit configured to generate second execution data by operating on the first register and the second register and to store the second execution data in the second register, wherein the second execution data includes a normalized range value and a normalized offset value.

13. The digital processor of claim 12, wherein the first register includes data fields for storing the current output bit, the temporary range value, the temporary offset value, a binary most-probable-symbol (MPS) value, a current probability state, two next probability state values, and a set of least-probable-symbol (LPS) range values.

14. The digital processor of claim 13, wherein the first register also includes data fields for storing a bypass control bit, a termination control bit and a zero state indicator, wherein the zero state indicator indicates whether a current context has a probability state index equal to zero.

15. The digital processor of claim 13, wherein first execution data also include an update for the current probability state and an update for the binary MPS value, wherein the first execution unit is configured to initiate a memory access and to store results of the memory access in the

first register, wherein the results of the memory access include updates for the two next probability state values and for the set of LPS range values.

16. The digital processor of claim 12 further comprising additional execution units configured to execute instructions of a baseline instruction set, wherein the baseline instruction set includes a set of integer arithmetic operations and a set of floating point arithmetic operations.

17. The digital processor of claim 12, wherein the second register includes data fields for storing a first stream buffer, the normalized range value and the normalized offset value, wherein the first stream buffer stores bits of the encoded bit stream.

18. The digital processor of claim 17, wherein the second execution unit is configured to access a number of bits from the first stream buffer, wherein the normalized offset value includes one or more of said number of accessed bits.

19. The digital processor of claim 17 wherein second register includes a remaining bit count that indicates a number of valid bits remaining in the first stream buffer.

20. A method for performing CABAC decoding, the method comprising:
executing a first instruction on a first execution unit of a microprocessor, wherein said executing the first instruction includes generating a current output bit, a temporary range value and a temporary offset value based at least on a current range value and a current offset value, wherein the current offset value includes one or more bits of the input bitstream; and
executing a second instruction on a second execution unit of the microprocessor, wherein said executing the second instruction includes generating an update for the current range value and an update for the current offset value by normalizing respectively the temporary range value and the temporary offset value.




FIG. 1




FIG. 2A




FIG. 2B




FIG. 3




FIG. 4A





FIG. 4B



6 / 18




FIG. 6




FIG. 7




FIG. 8




FIG. 9




FIG. 10

9 / 18




FIG. 11




FIG. 12




FIG. 13

11 / 18






FIG. 14

12 / 18








FIG. 17




FIG. 18

16 / 18

```

unsigned short RangeLPS, codlRange;
unsigned short tmpRange, x;
unsigned short tmpOffsetA, tmpOffsetB;
unsigned short tmpState, LZCount;
bit RangeIDX[1:0], tmpQ, binVal, mode, modeS;

// intermediate result calculation

RangeIDX = src2.Range >> 6; // upper two bits of Range
RangeLPS = src1.rLPS[RangeIDX]; // select one of the four rLPS values
if (src1.ByPass) x = 0
else
    if (src1.Terminate) x = 2
    else x = RangeLPS;
codlRange = src2.Range - x;
tmpOffsetA = src1.ByPass ? ((src2.Offset << 1) || src2.SBuf1[0]) : (src2.Offset);
tmpOffsetB = tmpOffsetA - codlRange;
tmpQ = tmpOffsetB.Carry; // carry from calculating tmpOffsetB
dest.NextOffset = tmpQ ? tmpOffsetA : tmpOffsetB;

mode = tmpQ || src1.Terminate || src1.ByPass;
tmpRange = mode ? codlRange : RangeLPS;
dest.NextRange = tmpRange;

modeS = src1.Terminate || src1.ByPass;
dest.MPS = ((src1.StateZ && ~tmpQ) ^ (src1.MPS));

tmpState = tmpQ ? src1.nextStateMPS : src1.nextStateLPS;
dest.CurrentState = tmpState;
dest.StateZ = (tmpState == 0);
dest.MPSprime = src1.MPS;
binVal = modeS ? (~tmpQ) : (~tmpQ ^ src1.MPS);
dest.Q = binVal;
dest.Symbol = ((src1.Symbol << 1) || binVal);

LZCount = LeadingZeroCount(tmpRange); // count the number of leading zeros
dest.SCount = modeS ? (tmpQ && src1.Terminate) || src1.ByPass : LZCount;

// update context state for the current context
// mem64 is a valid 64-bit aligned memory address,

MemLoad = Memory[mem64 + (tmpState * 8)];
dest.nextStateLPS = MemLoad.nextStateLPS;
dest.nextStateMPS = MemLoad.nextStateMPS;
dest.Terminate = MemLoad.Terminate;
dest.ByPass = MemLoad.ByPass;
dest.rLPS[3:0] = MemLoad.rLPS[3:0];

```

FIG. 19

17 / 18

```
unsigned short tmpBitsLeft, tmpLcnt, tmpAdd;
unsigned long tmpVal, tmpBuf, tmpBuf2;
bit tmpGet;

// intermediate result calculation

tmpBitsLeft = src1.BitsLeft - src2.Scount;
tmpGet = (tmpBitsLeft < 1);
dest.GetBits = tmpGet; // may be used to set flag in GPR version

tmpAdd = tmpGet ? 32 : 0;
dest.BitsLeft = tmpBitsLeft + tmpAdd;

dest.Range = src2.NextRange << src2.SCount;
dest.ShCnt = src2.Scount;

tmpLcnt = src2.Scount - src1.BitsLeft; // yes, this is -tmpBitsLeft
tmpLS = tmpGet ? tmpLcnt : 0;
tmpBuf = src1.buf2 << tmpLS;
tmpBuf2 = tmpGet ? (tmpBuf) : (0);

dest.SBuf2 = tmpGet ? tmpBuf2[31:0] : 0;

tmpVal = (bitreverse(src2.NextOffset) || (src1.SBuf1 << 9) >> src2.SCount;
dest.SBuf1 = (tmpVal + tmpBuf2) >> 9;
dest.Offset = (tmpVal + tmpBuf2) & 0xFF;
```

FIG. 20

18 / 18




FIG. 21

INTERNATIONAL SEARCH REPORT

International application No

PCT/US2010/033744

A. CLASSIFICATION OF SUBJECT MATTER

INV. H04N7/26

ADD.

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

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

H04N

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

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

EPO-Internal, COMPENDEX, 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	JVT: "Draft ITU-T Recommendation and Final Draft International Standard of Joint Video Specification (ITU-T Rec. H.264 ISO/IEC 14496-10 AVC)" ITU STUDY GROUP 16 - VIDEO CODING EXPERTS GROUP -ISO/IEC MPEG & ITU-T VCEG(ISO/IEC JTC1/SC29/WG11 AND ITU-T SG16 Q6), XX, XX, no. JVT-G050r1, 14 March 2003 (2003-03-14), XP030005712 cited in the application page 187 - page 193, paragraph 9.3.3.2; figures 9-2, 9-3, 9-5, 9-6 page 193 - page 198, paragraph 9.3.4 page 171, paragraph 9.3.1.2 ----- -/--	1-20



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 document 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

27 August 2010

Date of mailing of the international search report

03/09/2010

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

Martinière, Anthony

INTERNATIONAL SEARCH REPORT

International application No

PCT/US2010/033744

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2006/294174 A1 (HAQUE MUNSI A [US] ET AL) 28 December 2006 (2006-12-28) paragraph [0042] - paragraph [0059]; figures 3,5 paragraph [0060] - paragraph [0068]; figure 7 -----	1-20
X	MARPE D ET AL: "Context-based adaptive binary arithmetic coding in the H.264/AVC video compression standard" IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS FOR VIDEO TECHNOLOGY, IEEE SERVICE CENTER, PISCATAWAY, NJ, US LNKD- DOI:10.1109/TCSVT.2003.815173, vol. 13, no. 7, 1 July 2003 (2003-07-01), pages 620-636, XP011099255 ISSN: 1051-8215 * abstract page 621 - page 627, paragraph II page 627 - page 634, paragraph III -----	1-20
X	US 2007/297501 A1 (HUSSAIN ZAHID [GB] ET AL HUSSAIN ZAHID [GB] ET AL) 27 December 2007 (2007-12-27) paragraph [0034] - paragraph [0055]; figures 3,4 paragraph [0056] - paragraph [0068]; figure 5A paragraph [0069] - paragraph [0074]; figure 5B paragraph [0075] - paragraph [0121]; figures 5C, 6A -----	1-20
A	US 5 909 572 A (THAYER JOHN S [US] ET AL) 1 June 1999 (1999-06-01) column 14; figure 7 -----	1-19
A	DASU A ET AL: "A Survey of Media Processing Approaches" IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS FOR VIDEO TECHNOLOGY, IEEE SERVICE CENTER, PISCATAWAY, NJ, US, vol. 12, no. 8, 1 August 2002 (2002-08-01) , XP011071857 ISSN: 1051-8215 the whole document -----	1-19
A,P	US 2009/160863 A1 (FRANK MICHAEL [US]) 25 June 2009 (2009-06-25) * abstract paragraph [0024] - paragraph [0045]; figure 1 -----	1-19

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2010/033744

Patent document cited in search report		Publication date	Patent family member(s)		Publication date
US 2006294174	A1	28-12-2006	NONE		
US 2007297501	A1	27-12-2007	NONE		
US 5909572	A	01-06-1999	US	6298438 B1	02-10-2001
			US	6141673 A	31-10-2000
US 2009160863	A1	25-06-2009	WO	2009082428 A1	02-07-2009