



(51) International Patent Classification:

G06F 9/38 (2018.01) H03M 7/30 (2006.01)

G06F 9/30 (2018.01)

(21) International Application Number:

PCT/EP2020/053383

(22) International Filing Date:

11 February 2020 (11.02.2020)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

16/286,905 27 February 2019 (27.02.2019) US

(71) Applicant: **INTERNATIONAL BUSINESS MACHINES CORPORATION** [US/US]; New Orchard Road, Armonk, New York 10504 (US).

(71) Applicant (for MG only): **IBM DEUTSCHLAND GMBH** [DE/DE]; c/o IBM Deutschland Management & Business

Support GmbH, Patentwesen und Urheberrecht, IBM Allee 1, 71139 Ehningen (DE).

(72) Inventors: **WEISHAUPT, Simon**; c/o IBM Deutschland Research & Development GmbH, Schoenaicher Strasse 220, 71032 Boeblingen (DE). **KLEIN, Matthias**; c/o IBM Corp., 2455 South Road, Poughkeepsie, New York 12601 (US). **USENBINZ, Stefan**; c/o IBM Deutschland Research & Development GmbH, Schoenaicher Strasse 220, 71032 Boeblingen (DE). **SOFIA, Anthony**; c/o IBM Corp., 2455 South Road, Poughkeepsie, New York 12601 (US).

(74) Agent: **DOEHLER, Denis**; c/o IBM Deutschland Management & Business Support GmbH, Patentwesen und Urheberrecht, IBM Allee 1, 71139 Ehningen (DE).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN,

(54) Title: SECURE AND EFFICIENT APPLICATION DATA PROCESSING




FIG. 22

(57) Abstract: A system architecture is provided and includes first and second processing units respectively communicative with an on-chip coherency unit and an accelerator communicative with the on-chip coherency unit. The accelerator is configured to execute an operation responsive to a call issued by one of the first and second processing units. The first processing unit is configured to set an asynchronous operation flag (AOF) to indicate that the second processing unit is to conduct an operation for the first processing unit. The second processing unit is configured to respond to the AOF by building scatter gather lists and subsequently issuing the call and feeding the scatter gather lists to the accelerator to facilitate execution of the operation by the accelerator.

HR, HU, ID, IL, IN, IR, IS, JO, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, WS, ZA, ZM, ZW.

**(84) Designated States** (*unless otherwise indicated, for every kind of regional protection available*): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

**Published:**

- *with international search report (Art. 21(3))*
- *in black and white; the international application as filed contained color or greyscale and is available for download from PATENTSCOPE*

## SECURE AND EFFICIENT APPLICATION DATA PROCESSING

## BACKGROUND

[0001] The present invention generally relates to data processing, and more specifically, to secure and efficient application data processing.

5 [0002] In computing, asynchronous execution of an operation, such as a DEFLATE compression call (DFLTCC), occurs when a processor that does not want to wait for the execution of the DFLTCC asks another (assist) processor to execute the DFLTCC on its behalf. The assist processor then runs firmware (e.g., i390 firmware) in order to execute the DFLTCC and reports back to the requesting processor.

10 SUMMARY

[0003] In accordance with embodiments of the present invention, a system architecture is provided and includes first and second processing units respectively communicative with an on-chip coherency unit and an accelerator communicative with the on-chip coherency unit. The accelerator is configured to execute an operation responsive to a  
15 call issued by one of the first and second processing units. The first processing unit is configured to set an asynchronous operation flag (AOF) to indicate that the second processing unit is to conduct an operation for the first processing unit. The second processing unit is configured to respond to the AOF by building scatter gather lists and subsequently issuing the call and feeding the scatter gather lists to the accelerator to facilitate execution of the  
20 operation by the accelerator.

[0004] In accordance with additional or alternative embodiments of the present invention, the first and second processing units are each respectively provided as ones of multiple processing units respectively communicative with the on-chip coherency unit via first interfaces.

25 [0005] In accordance with additional or alternative embodiments of the present invention, the on-chip coherency unit is communicative with external processors and controllers.

[0006] In accordance with additional or alternative embodiments of the present invention, the system architecture further includes page buffer units (PBUs) respectively communicative with the on-chip coherency unit via second interfaces and a memory core.

5 [0007] In accordance with additional or alternative embodiments of the present invention, the accelerator is communicative with the on-chip coherency unit via a second interface comprising a direct memory access (DMA)-like interface.

[0008] In accordance with additional or alternative embodiments of the present invention, the operation includes a DEFLATE operation.

10 [0009] In accordance with additional or alternative embodiments of the present invention, the first processing unit is further configured to determine whether to operate in one of a synchronous mode or an asynchronous mode, issue the call in an event the first processing unit determines that the synchronous mode should be in effect and set the AOF in an event the first processing unit determines that the asynchronous mode should be in effect.

15 [0010] In accordance with additional or alternative embodiments of the present invention, the scatter gather lists include absolute addresses and associated lengths for source data and target data of the data compression.

[0011] In accordance with additional or alternative embodiments of the present invention, the absolute addresses of the source data are non-consecutive in absolute address space and the absolute addresses of the target data are non-consecutive in absolute address  
20 space.

[0012] In accordance with additional or alternative embodiments of the present invention, the absolute addresses of the source data are non-sequential and the absolute addresses of the target data are non-sequential.

25 [0013] In accordance with embodiments of the present invention, a system architecture is provided and includes one or more drawers, each of which includes one or more central processors (CPs) and one or more system controllers (SCs). Each of the one or more CPs includes an accelerator configured to execute data compression faster than the first

and second processing units responsive to a call issued by one of the first and second processing units. The first processing unit is configured to set an asynchronous compression flag (ACF) to indicate that the second processing unit is to conduct a compression operation for the first processing unit. The second processing unit is configured to respond to the ACF by building scatter gather lists and subsequently issuing the call and feeding the scatter gather lists to the accelerator to facilitate execution of the data compression by the accelerator.

[0014] In accordance with additional or alternative embodiments of the present invention, the one or more drawers are communicative with each other and the on-chip coherency unit of each of the one or more CPs of each of the one or more drawers is communicative with the one or more CPs and the one or more SCs of the corresponding drawer.

[0015] In accordance with additional or alternative embodiments of the present invention, the first and second processing units of each of the one or more CPs of each of the one or more drawers are provided as ones of multiple processing units respectively communicative with a corresponding on-chip coherency unit via first interfaces.

[0016] In accordance with additional or alternative embodiments of the present invention, each of the one or more CPs of each of the one or more drawers further includes page buffer units (PBUs) respectively communicative with a corresponding on-chip coherency unit via second interfaces and a memory core.

[0017] In accordance with additional or alternative embodiments of the present invention, the accelerator of each of the one or more CPs of each of the one or more drawers is communicative with a corresponding on-chip coherency unit via a second interface including a direct memory access (DMA)-like interface.

[0018] In accordance with additional or alternative embodiments of the present invention, the compression operation includes a DEFLATE operation.

[0019] In accordance with additional or alternative embodiments of the present invention, the first processing unit is further configured to determine whether to operate in one of a synchronous mode or an asynchronous mode, issue the call in an event the first

processing unit determines that the synchronous mode should be in effect and set the ACF in an event the first processing unit determines that the asynchronous mode should be in effect.

[0020] In accordance with additional or alternative embodiments of the present invention, the scatter gather lists include absolute addresses and associated lengths for source data and target data of the operation.

[0021] In accordance with additional or alternative embodiments of the present invention, the absolute addresses of the source data are non-consecutive and the absolute addresses of the target data are non-consecutive.

[0022] In accordance with additional or alternative embodiments of the present invention, the absolute addresses of the source data are non-sequential in absolute address space and the absolute addresses of the target data are non-sequential in absolute address space.

[0023] In accordance with embodiments of the present invention, a computer program product is provided. The computer program product includes a processing circuit and a memory unit. The memory unit has executable instructions stored thereon, which, when executed cause the processing circuit to execute a method of securely and efficiently executing data compression or decompression. The method includes determining whether to operate in one of a synchronous mode or an asynchronous mode, setting an asynchronous operation flag (AOF) in an event the determining indicates that the asynchronous mode should be in effect and building scatter gather lists to be fed to an accelerator to facilitate execution of the data compression or decompression by the accelerator in response to the setting of the AOF.

[0024] In accordance with additional or alternative embodiments of the present invention, the scatter gather lists include absolute addresses and associated lengths for source data and target data of the data compression or decompression and the absolute addresses of the source data are non-consecutive and the absolute addresses of the target data are non-consecutive.

[0025] In accordance with embodiments of the present invention, a computer-implemented method of securely and efficiently executing data compression or decompression is provided. The computer-implemented method includes determining whether to operate in one of a synchronous mode or an asynchronous mode, setting an asynchronous operation flag (AOF) in an event the determining indicates that the asynchronous mode should be in effect and building scatter gather lists to be fed to an accelerator to facilitate execution of the data compression or decompression by the accelerator in response to the setting of the AOF.

[0026] In accordance with additional or alternative embodiments of the present invention, the scatter gather lists include absolute addresses and associated lengths for source data and target data of the data compression or decompression and the absolute addresses of the source data are non-consecutive and the absolute addresses of the target data are non-consecutive.

[0027] In accordance with embodiments of the present invention, a system architecture is provided. The system architecture includes first and second processing units respectively communicative with an on-chip coherency unit and an accelerator communicative with the on-chip coherency unit. The accelerator is configured to execute data compression or decompression faster than the first and second processing units responsive to a call issued by one of the first and second processing units. The first processing unit is configured to execute a method including determining whether to operate in one of a synchronous mode or an asynchronous mode and setting an asynchronous operation flag (AOF) in an event the determining indicates that the asynchronous mode should be in effect. The second processing unit is configured to execute a method including building scatter gather lists to be fed to the accelerator to facilitate execution of the data compression or decompression by the accelerator in response to the setting of the AOF.

[0028] Technical effects and benefits of the embodiments described herein provide for executions of certain operations in an asynchronous mode in which the ACF is set and millicode thus takes absolute addresses from scatter gather lists and feeds those addresses to an accelerator.

[0029] Additional technical features and benefits are realized through the techniques of the present invention. Embodiments and aspects of the invention are described in detail herein and are considered a part of the claimed subject matter. For a better understanding, refer to the detailed description and to the drawings.

5

## BRIEF DESCRIPTION OF THE DRAWINGS

[0030] The specifics of the exclusive rights described herein are particularly pointed out and distinctly claimed in the claims at the conclusion of the specification. The foregoing and other features and advantages of the embodiments of the invention are apparent from the following detailed description taken in conjunction with the accompanying drawings in  
10 which:

[0031] FIG. 1A depicts one example of a computing environment to incorporate and use one or more aspects of the present invention;

[0032] FIG. 1B depicts further details of a processor of FIG. 1A, in accordance with one or more aspects of the present invention;

15 [0033] FIG. 2 depicts another example of a computing environment to incorporate and use one or more aspects of the present invention;

[0034] FIG. 3A depicts one format of a DEFLATE Conversion Call (DFLTCC) instruction, in accordance with an aspect of the present invention;

[0035] FIG. 3B depicts one example of fields of an implied register, general register  
20 0, used by the DEFLATE Conversion Call instruction, in accordance with an aspect of the present invention;

[0036] FIG. 3C depicts one example of function codes for the DEFLATE Conversion Call instruction, in accordance with an aspect of the present invention;

[0037] FIG. 3D depicts one example of a field of an implied register, general register  
25 1, used by the DEFLATE Conversion Call instruction, in accordance with an aspect of the present invention;

[0038] FIG. 3E depicts one example of contents of a register, R1, specified by the DEFLATE Conversion Call instruction, in accordance with an aspect of the present invention;

[0039] FIG. 3F depicts one example of contents of a register, R1+1, used by the DEFLATE Conversion Call instruction, in accordance with an aspect of the present invention;

5 [0040] FIG. 3G depicts one example of contents of a register, R2, specified by the DEFLATE Conversion Call instruction, in accordance with an aspect of the present invention;

[0041] FIG. 3H depicts one example of contents of a register, R2+1, used by the DEFLATE Conversion Call instruction, in accordance with an aspect of the present invention;

[0042] FIG. 3I depicts one example of contents of a register, R3, specified by the  
10 DEFLATE Conversion Call instruction, in accordance with an aspect of the present invention;

[0043] FIG. 3J depicts one example of contents of a parameter block used by the DFLTCC-QAF (query available functions) function of the DEFLATE Conversion Call instruction, in accordance with an aspect of the present invention;

[0044] FIG. 3K depicts one example of contents of a parameter block used by the  
15 DFLTCC-GDHT (generate dynamic-Huffman table) function of the DEFLATE Conversion Call instruction, in accordance with an aspect of the present invention;

[0045] FIG. 3L depicts one example of contents of a parameter block used by the DFLTCC-CMPR (compress) and DFLTCC-XPND (expand) functions of the DEFLATE Conversion Call instruction, in accordance with an aspect of the present invention;

20 [0046] FIG. 4 depicts one example of a sub-byte boundary, in accordance with one or more aspects of the present invention;

[0047] FIGS. 5A-5C depict examples illustrating how a sub-byte boundary applies to the DFLTCC-CMPR function, in accordance with an aspect of the present invention;

[0048] FIG. 6 depicts one example of a block of data with no compression, in  
25 accordance with an aspect of the present invention;

[0049] FIG. 7 depicts one example of a block with compressed data using a fixed-Huffman table (FHT), in accordance with an aspect of the present invention;

[0050] FIG. 8 depicts one example of a block with compressed data using a dynamic-Huffman table (DHT), in accordance with an aspect of the present invention;

5 [0051] FIG. 9 depicts one example of a compressed data set in storage, in accordance with an aspect of the present invention;

[0052] FIG. 10 depicts one example of a sample of a program compressing data into three blocks of a compressed data set, in accordance with an aspect of the present invention;

10 [0053] FIG. 11 depicts one example of parameter block contents for a DFLTCC-CMPR function operating on a first compressed data block of a set, in accordance with an aspect of the present invention;

[0054] FIG. 12 depicts one example of parameter block contents for a DFLTCC-CMPR function operating on a second compressed data block of a set, in accordance with an aspect of the present invention;

15 [0055] FIG. 13 depicts one example of a sample of a program decompressing data from a compressed data set, in accordance with an aspect of the present invention;

[0056] FIGS. 14A-14C depict examples of an in-line history buffer before and after executing DFLTCC-CMPR multiple times, in accordance with an aspect of the present invention;

20 [0057] FIGS. 15A-15E depict examples of a circular history buffer before and after executing DFLTCC multiple times, in accordance with an aspect of the present invention;

[0058] FIGS. 16A-16C depict examples of an in-line history buffer before and after executing DFLTCC-XPND multiple times, in accordance with an aspect of the present invention;

[0059] FIG. 17 depicts one example of using the DEFLATE Conversion Call instruction, in accordance with an aspect of the present invention;

[0060] FIG. 18 depicts one example of using a circular history buffer, in accordance with an aspect of the present invention

5 [0061] FIG. 19 depicts a system architecture in accordance with embodiments of the present invention;

[0062] FIG. 20 depicts additional features of the system architecture of FIG. 19 in accordance with embodiments of the present invention;

10 [0063] FIG. 21 depicts additional features of the system architecture of FIG. 19 in accordance with embodiments of the present invention;

[0064] FIG. 22 is a flow diagram illustrating a method or a computer-implemented method of securely and efficiently executing data compression in accordance with embodiments of the present invention;

15 [0065] FIG. 23 is a schematic diagram illustrating relative positioning of source and target data of a data compression operation in absolute memory space in accordance with embodiments of the present invention; and

[0066] FIG. 24 is a schematic diagram illustrating scatter gather lists generated for handling the source and target data of FIG. 23 in accordance with embodiments of the present invention.

20 [0067] The diagrams depicted herein are illustrative. There can be many variations to the diagram or the operations described therein without departing from the spirit of the invention. For instance, the actions can be performed in a differing order or actions can be added, deleted or modified. Also, the term “coupled” and variations thereof describes having a communications path between two elements and does not imply a direct connection between  
25 the elements with no intervening elements/connections between them. All of these variations are considered a part of the specification.

[0068] In the accompanying figures and following detailed description of the disclosed embodiments, the various elements illustrated in the figures are provided with two or three digit reference numbers. With minor exceptions, the leftmost digit(s) of each reference number correspond to the figure in which its element is first illustrated.

5

## DETAILED DESCRIPTION

[0069] Various embodiments of the invention are described herein with reference to the related drawings. Alternative embodiments of the invention can be devised without departing from the scope of this invention. Various connections and positional relationships (e.g., over, below, adjacent, etc.) are set forth between elements in the following description and in the drawings. These connections and/or positional relationships, unless specified otherwise, can be direct or indirect, and the present invention is not intended to be limiting in this respect. Accordingly, a coupling of entities can refer to either a direct or an indirect coupling, and a positional relationship between entities can be a direct or indirect positional relationship. Moreover, the various tasks and process steps described herein can be incorporated into a more comprehensive procedure or process having additional steps or functionality not described in detail herein.

[0070] The following definitions and abbreviations are to be used for the interpretation of the claims and the specification. As used herein, the terms “comprises,” “comprising,” “includes,” “including,” “has,” “having,” “contains” or “containing,” or any other variation thereof, are intended to cover a non-exclusive inclusion. For example, a composition, a mixture, process, method, article, or apparatus that comprises a list of elements is not necessarily limited to only those elements but can include other elements not expressly listed or inherent to such composition, mixture, process, method, article, or apparatus.

25 [0071] Additionally, the term “exemplary” is used herein to mean “serving as an example, instance or illustration.” Any embodiment or design described herein as “exemplary” is not necessarily to be construed as preferred or advantageous over other embodiments or designs. The terms “at least one” and “one or more” may be understood to include any integer number greater than or equal to one, i.e. one, two, three, four, etc. The

terms “a plurality” may be understood to include any integer number greater than or equal to two, i.e. two, three, four, five, etc. The term “connection” may include both an indirect “connection” and a direct “connection.”

[0072] The terms “about,” “substantially,” “approximately,” and variations thereof, are intended to include the degree of error associated with measurement of the particular quantity based upon the equipment available at the time of filing the application. For example, “about” can include a range of  $\pm 8\%$  or  $5\%$ , or  $2\%$  of a given value.

[0073] For the sake of brevity, conventional techniques related to making and using aspects of the invention may or may not be described in detail herein. In particular, various aspects of computing systems and specific computer programs to implement the various technical features described herein are well known. Accordingly, in the interest of brevity, many conventional implementation details are only mentioned briefly herein or are omitted entirely without providing the well-known system and/or process details.

[0074] Turning now to an overview of technologies that are more specifically relevant to aspects of the invention, an application running on a core processing unit (PU) operates in logical address space where each logical address that the processor sees in logical address space (which is private to the application) maps to an address in absolute address space (which is the common system wide address). If that processor executes the application itself, the execution is synchronous. Here, the source data is typically located in consecutive areas in logical address space and the target data is also typically located in consecutive areas in logical address space. On the other hand, in an event the processor does not want to wait for an execution of the application, the processor can ask another (assist) processor to execute the application on its behalf. The assist processor then runs firmware in order to execute the application in an asynchronous execution and reports back to the requesting processor upon completion. In doing so, the assist processor is unaware of the mapping between logical and absolute addresses and, therefore, the assist processor needs to operate directly on absolute addresses. Here, the source and target data are typically not located in consecutive areas in absolute address space.

[0075] In asynchronous executions, such as for asynchronous DEFLATE compression calls (DFLTCCs) or decompression calls, the assist processor has to copy data from non-consecutive source and target buffers and then place that copied data in hardware system areas (HSAs). This places a high demand on the system as additional, private areas in HSA  
5 are needed. There is also significant overhead due to copying and concatenating source data which reduces throughput and increases utilization. In addition, security concerns are presented due to customer data being placed in HSA.

[0076] Other issues with asynchronous executions are that the assist processor is often required to invoke certain instructions multiple times with each consecutive source and target  
10 data chunk being copied and concatenated. With each call having a relatively small consecutive chunk of source and target data, there can be a relatively high amount of instruction overhead and greatly reduced throughput.

[0077] Turning now to an overview of the aspects of the invention, one or more embodiments of the invention address the above-described shortcomings of the prior art by  
15 providing for an introduction of an asynchronous operation flag (AOF) or an asynchronous compression flag (ACF) for certain compression or decompression operations, such as DEFLATE compression calls (DFLTCC). The AOF or ACF indicates that an asynchronous mode is in effect and is set when an assist processor executes the certain operations on behalf of a different processor. For such asynchronous execution, firmware on the assist processor  
20 builds a source scatter gather list with system absolute addresses and associated lengths for source data and target data and calls the instruction with the AOF or the ACF and the following parameters: address of source scatter gather list, length of total source data, address of target scatter gather list and length of total target data. If the AOF or the ACF is set, millicode of the certain operations works with absolute addresses provided in the scatter  
25 gather lists instead of translating consecutive source and target buffers.

[0078] That is, for a DFLTCC in synchronous mode in which the ACF is not set, the DFLTCC millicode takes base offset of source and target buffers and translates each page (4k increments) from logical to absolute and feeds those addresses to an accelerator. By contrast, for the DFLTCC in asynchronous mode in which the ACF is set, the DFTLCC millicode

takes absolute addresses from the scatter gather list and feeds those addresses to the accelerator.

[0079] The above-described aspects of the invention address the shortcomings of the prior art by providing for executions of certain operations in an asynchronous mode in which the ACF is set and millicode thus takes absolute addresses from the scatter gather list and feeds those addresses to the accelerator.

[0080] Turning now to a more detailed description of aspects of the present invention, a capability is provided to facilitate processing within a computing environment. As one example, a single instruction (e.g., a single architected hardware machine instruction at the hardware/software interface) is provided to perform a function (also referred to as an operation), such as a compression or decompression function, to compress and/or decompress (also referred to as uncompress) data. The instruction is part of a general-purpose processor instruction set architecture (ISA), which is dispatched by a program (e.g., an operating system or a user program) on the general-purpose processor. By using an ISA instruction to perform compression/decompression, task switches by the operating system are not needed to perform the compression/decompression operations, thereby saving execution cycles. Further, by using a single instruction to compress and/or decompress data, execution time within a processor, such as a general-purpose processor, is reduced.

[0081] In one example, the instruction performs compressing and uncompressing operations compliant with an industry standard, referred to as the DEFLATE standard, and the instruction is referred to as a DEFLATE Conversion Call instruction. The DEFLATE standard includes descriptions for compressed data symbols which represent duplicate strings in the original form of the data (in the uncompressed form of the data). Such symbols include a pointer and a length of a duplicate string which describe the location and length of the duplicate string, which was previously processed, in relationship to the current location of data being processed. The previously processed uncompressed form of the data is referred to as history. In one example, the history is a contiguous number of bytes in memory, which can be as large as, for instance, 32 K-bytes.

[0082] One embodiment of a computing environment to incorporate and use one or more aspects of the present invention is described with reference to FIG. 1A. A computing environment 100 includes, for instance, a processor 102 (e.g., a central processing unit), a memory 104 (e.g., main memory; a.k.a., system memory, main storage, central storage, storage), and one or more input/output (I/O) devices and/or interfaces 106 coupled to one another via, for example, one or more buses 108 and/or other connections.

[0083] In one example, processor 102 is based on the z/Architecture® hardware architecture offered by International Business Machines Corporation, Armonk, New York, and is part of a server, such as an IBM Z® server, which is also offered by International Business Machines Corporation and implements the z/Architecture hardware architecture. One embodiment of the z/Architecture hardware architecture is described in a publication entitled, “z/Architecture Principles of Operation,” IBM Publication No. SA22-7832-11, 12th edition, September 2017, which is hereby incorporated herein by reference in its entirety. The z/Architecture hardware architecture, however, is only one example architecture; other architectures and/or other types of computing environments may include and/or use one or more aspects of the present invention. In one example, the processor executes an operating system, such as the z/OS® operating system, also offered by International Business Machines Corporation.

[0084] Processor 102 includes a plurality of functional components used to execute instructions. As depicted in FIG. 1B, these functional components include, for instance, an instruction fetch component 120 to fetch instructions to be executed; an instruction decode unit 122 to decode the fetched instructions and to obtain operands of the decoded instructions; an instruction execute component 124 to execute the decoded instructions; a memory access component 126 to access memory for instruction execution, if necessary; and a write back component 130 to provide the results of the executed instructions. One or more of these components may, in accordance with one or more aspects of the present invention, include at least a portion of or have access to one or more other components used in compress/decompress processing (or other processing that may use one or more aspects of the present invention), as described herein. The one or more other components include, for instance, a compress/decompress component (or other component) 136.

[0085] Another example of a computing environment to incorporate and use one or more aspects of the present invention is described with reference to FIG. 2. In one example, the computing environment is based on the z/Architecture hardware architecture; however, the computing environment may be based on other architectures offered by International Business Machines Corporation or others.

[0086] Referring to FIG. 2, in one example, the computing environment includes a central electronics complex (CEC) 200. CEC 200 includes a plurality of components, such as, for instance, a memory 202 (a.k.a., system memory, main memory, main storage, central storage, storage) coupled to one or more processors (a.k.a., central processing units (CPUs)) 204, and to an input/output subsystem 206.

[0087] Memory 202 includes, for example, one or more logical partitions 208, a hypervisor 210 that manages the logical partitions, and processor firmware 212. One example of hypervisor 210 is the Processor Resource/System Manager (PR/SM™) hypervisor, offered by International Business Machines Corporation, Armonk, New York. As used herein, firmware includes, e.g., the microcode of the processor. It includes, for instance, the hardware-level instructions and/or data structures used in implementation of higher level machine code. In one embodiment, it includes, for instance, proprietary code that is typically delivered as microcode that includes trusted software or microcode specific to the underlying hardware and controls operating system access to the system hardware.

[0088] Each logical partition 208 is capable of functioning as a separate system. That is, each logical partition can be independently reset, run a guest operating system 220 such as a z/OS operating system, or another operating system, and operate with different programs 222. An operating system or application program running in a logical partition appears to have access to a full and complete system, but in reality, only a portion of it is available.

[0089] Memory 202 is coupled to processors (e.g., CPUs) 204, which are physical processor resources that may be allocated to the logical partitions. For instance, a logical partition 208 includes one or more logical processors, each of which represents all or a share of a physical processor resource 204 that may be dynamically allocated to the logical partition.

[0090] Further, memory 202 is coupled to I/O subsystem 206. I/O subsystem 206 may be a part of the central electronics complex or separate therefrom. It directs the flow of information between main storage 202 and input/output control units 230 and input/output (I/O) devices 240 coupled to the central electronics complex.

5 [0091] Many types of I/O devices may be used. One particular type is a data storage device 250. Data storage device 250 may store one or more programs 252, one or more computer readable program instructions 254, and/or data, etc. The computer readable program instructions may be configured to carry out functions of embodiments of aspects of the invention.

10 [0092] As one example, each processor 204 includes at least one cache 260 (e.g., a local cache) of a cache hierarchy that includes a plurality of levels of cache, including one or more local caches and/or one or more shared caches. Further, in one embodiment, the local caches and memory 202 are coupled to a compress/decompress component (or other component) 262 used in performing one or more of compressing and/or decompressing of  
15 data (and/or other operations of one or more aspects of the present invention). In various examples, there may be one or more components performing these tasks. Many variations are possible.

[0093] In one embodiment, a processor (e.g., processor 204) obtains the instruction (e.g., the DEFLATE Conversion Call instruction), decodes the instruction, performs set-up  
20 for the instruction, including translating addresses to be used by the instruction, and sends commands for the instruction to a component coupled to the processor, such as component 262, to perform a function specified by the instruction. Component 262 has access to the cache hierarchy and memory, such that in performing the specified function, it reads data, processes it and stores the processed data back. As an example, component 262 is a hardware  
25 component.

[0094] In a further embodiment, at least a portion of component 262 is included as part of the processor. Many variations are possible.

[0095] Central electronics complex 200 may include and/or be coupled to removable/non-removable, volatile/non-volatile computer system storage media. For example, it may include and/or be coupled to a non-removable, non-volatile magnetic media (typically called a "hard drive"), a magnetic disk drive for reading from and writing to a removable, non-volatile magnetic disk (e.g., a "floppy disk"), and/or an optical disk drive for reading from or writing to a removable, non-volatile optical disk, such as a CD-ROM, DVD-ROM or other optical media. It should be understood that other hardware and/or software components could be used in conjunction with central electronics complex 200. Examples include, but are not limited to: microcode, device drivers, redundant processing units, external disk drive arrays, RAID systems, tape drives, and data archival storage systems, etc.

[0096] Further, central electronics complex 200 may be operational with numerous other general purpose or special purpose computing system environments or configurations. Examples of well-known computing systems, environments, and/or configurations that may be suitable for use with central electronics complex 200 include, but are not limited to, personal computer (PC) systems, server computer systems, thin clients, thick clients, handheld or laptop devices, multiprocessor systems, microprocessor-based systems, set top boxes, programmable consumer electronics, network PCs, minicomputer systems, mainframe computer systems, and distributed cloud computing environments that include any of the above systems or devices, and the like.

[0097] Although various examples of computing environments are described herein, one or more aspects of the present invention may be used with many types of environments. The computing environments provided herein are only examples.

[0098] In accordance with an aspect of the present invention, a computing environment, such as computing environment 100 or central electronics complex 200, employs a conversion facility that provides a mechanism to compress and decompress data. In one example, the conversion facility is a DEFLATE conversion facility that provides a mechanism to compress and decompress data using the DEFLATE compressed data format. In one example, the conversion facility is installed in the system when a facility indicator is set, e.g., to one. As one particular example of the z/Architecture hardware architecture, facility bit 151 is set to, e.g., one, when the conversion facility is installed in the

z/Architecture architectural mode. The facility includes, for instance, the DEFLATE Conversion Call instruction, an embodiment of which is described below.

[099] In one example, the DEFLATE Conversion Call instruction performs functions related to transforming the state of data between the original (uncompressed) form of the data, and a compressed representation of the data, as specified by a selected standard, such as the IETF (Internet Engineering Task Force) RFC (Request for Comments) 1951 specification, which is described in the DEFLATE Compressed Data Format Specification version 1.3 Internet Engineering Task Force, Request for Comments 1951, May 1996.

[0100] In one example, the uncompressed data is a sequence of bytes, and the compressed representation of the data includes symbols. Symbols represent an individual byte of uncompressed data, referred to as a literal byte, or represent a reoccurring sequence of bytes of uncompressed data, referred to as a duplicate string. A Huffman table, as an example, specifies the encoding and decoding between compressed data symbols and uncompressed data. There are two types of Huffman tables: a fixed-Huffman table (FHT), which is a predetermined specification which includes, for instance, all possible codings; and a dynamic-Huffman table (DHT), which is a set of codings created specifically for the data to be compressed, which may be a subset of all possible codings. A compressed representation of data generated with a DHT is typically smaller than a compressed representation of the same data generated with an FHT. A portion of the most recently processed uncompressed data, referred to as history, is maintained for encoding and decoding compressed data symbols representing duplicate strings. The history is the reference source for duplicate strings. The history is updated as data is processed during an operation.

[0101] As indicated, in one example, the DEFLATE Conversion Call instruction uses the DEFLATE compressed data format, which is described in RCF 1951, DEFLATE Compressed Data Format Specification version 1.3. Attributes of the DEFLATE standard which apply to the DEFLATE Conversion Call instruction include, for instance:

[0102] • A compressed data set includes a series of blocks. There are three types of blocks. One type includes a 3-bit header followed by length information and uncompressed data, and two types of blocks include a 3-bit header followed by compressed data elements.

[0103] • Compressed data elements may include a compressed representation of a dynamic-Huffman table, compressed data symbols, and an end-of-block (EOB) symbol.

[0104] • Compressed data elements have various bit lengths.

[0105] • Compressed data elements may begin or end between byte boundaries  
5 in storage.

[0106] • Compressed data elements are loaded into bytes in order from, e.g., the rightmost bit position to the leftmost bit position.

[0107] When a compressed data element occupies part of, and not all of, a byte in storage, the entire byte in storage is accessed. Storage operand lengths specify the number of  
10 addressable bytes, which may specify more bits than the compressed data occupies.

[0108] Additional details regarding compressed data blocks are described further below.

[0109] One embodiment of the DEFLATE Conversion Call (DFLTCC) instruction is described with reference to FIGS. 3A-3L. The instruction is executed, in one example, using  
15 a general-purpose processor (e.g., processor 102 or 204). In the description herein, specific locations, specific fields and/or specific sizes of the fields are indicated (e.g., specific bytes and/or bits). However, other locations, fields and/or sizes may be provided. Further, although the setting of a bit to a particular value, e.g., one or zero, is specified, this is only an example. The bit may be set to a different value, such as the opposite value or to another  
20 value, in other examples. Many variations are possible.

[0110] In one embodiment, a program (e.g., an operating system or user program) may execute the DEFLATE Conversion Call instruction multiple times to compress or uncompress a single data stream. For instance, when an application compresses or decompresses a large data stream (e.g., greater than 1 M-bytes), the operation may include  
25 multiple calls to compress or decompress buffered portions of the data stream. In accordance with one aspect of the present invention, a program declares a buffer (e.g., a 32 K-byte buffer), which is used to accumulate the history of uncompressed data processed during an

operation which spans multiple executions of the DEFLATE Conversion Call instruction. The buffer is referred to as a circular history buffer, which is defined using the DEFLATE Conversion Call instruction, as described herein.

[0111] Referring to FIG. 3A, in one example, a format of a DEFLATE Conversion

5 Call (DFLTCC) instruction 300 is an RRF format that denotes a register and register operation with an extended operation code (opcode) field and an additional register field. As an example, the instruction includes an operation code field 302 (e.g., bits 0-15) having an operation code indicating a DEFLATE Conversion Call operation; a first register field (R1) 304 (e.g., bits 24-27) designating a first pair of general registers; a second register field (R2) 10 306 (e.g., bits 28-31) designating a second pair of general registers; and a third register field (R3) 308 (e.g., bits 16-19) designating a third general register. The contents of a register designated by R1 field 304 specify a location of a first operand (in storage); the contents of a register designated by R2 field 306 specify a location of a second operand (in storage); and the contents of a register designated by R3 field 308 specify a location of a third operand (in 15 storage). The contents of R1 + 1 specify the length of the first operand, and the contents of R2 + 1 specify the length of the second operand. In one example, bits 20-23 of the instruction are reserved and should contain zeros; otherwise, the program may not operate compatibly in the future. As used herein, the program is the one issuing the DEFLATE Conversion Call instruction. It may be a user program, an operating system or another type of program.

20 [0112] In one embodiment, execution of the instruction includes the use of one or more implied general registers (i.e., registers not explicitly designated by the instruction). For instance, general registers 0 and 1 are used in execution of the DEFLATE Conversion Call instruction, as described herein. General register 0 is used, in one example, to specify a function to be performed (as well as a history buffer type, described below), and general 25 register 1 is used to provide a location of a parameter block used by the instruction.

[0113] As an example, with reference to FIG. 3B, a general register 0 (309) includes a history buffer type field 310 and a function code field 312. In one particular example, bit position 56 of general register 0 includes a history buffer type, and bit positions 57-63 of general register 0 contain a function code; but in other embodiments, other bits may be used 30 to contain the history buffer type and/or the function code. When bits 57-63 of general

register 0 designate an unassigned or uninstalled function code, a specification exception is recognized, in one example.

[0114] Example assigned function codes for the DEFLATE Conversion Call instruction are shown in FIG. 3C and include, for instance: function code 0 (313) indicating a DFLTCC-QAF (query available functions) function; function code 1 (315) indicating a DFLTCC-GDHT (Generate Dynamic-Huffman Table) function; function code 2 (317) indicating a DFLTCC-CMPR (compress) function; and function code 4 (319) indicating a DFLTCC-XPND (expand) function. Each code uses a parameter block and the size of the parameter block depends, in one example, on the function. For instance, for the DFLTCC-QAF function, the parameter block is 32 bytes; for the DFLTCC-GDHT function, the parameter block is 384 bytes; and for the DFLTCC-CMPR and DFLTCC-XPND functions, the parameter block is 1536 bytes. Other function codes are unassigned in this example. Although example functions and function codes are described, other functions and/or function codes may be used.

[0115] When the specified function is DFLTCC-CMPR or DFLTCC-XPND, bit 56 of general register 0 specifies the history buffer type (HBT) used during the operation. When HBT is zero, the history buffer is called an in-line history buffer. When using an in-line history buffer, the history is, e.g., immediately to the left of the second operand when DFLTCC-CMPR is specified and is, e.g., immediately to the left of the first operand when DFLTCC-XPND is specified. When HBT is one, the history buffer is called a circular history buffer. When using a circular history buffer, the history is a portion of, or all of, the third operand when either DFLTCC-CMPR or DFLTCC-XPND is specified. When the DFLTCC-QAF or DFLTCC-GDHT function is specified, bit 56 of general register 0 is ignored. In one example, bit positions 0-31 of general register 0 are ignored. Further, in one example, bit positions 32-55 of general register 0 are reserved and should contain zeros; otherwise, the program may not operate compatibly in the future.

[0116] Further details regarding another implied register, general register 1, used by the DEFLATE Conversion Call instruction are described with reference to FIG. 3D. The contents of general register 1 (314) specify, for instance, a logical address 316 of the leftmost byte of a parameter block in storage. The parameter block is to be designated on a 4K-byte

boundary, in one example; otherwise a specification exception is recognized. Further details regarding the parameter block are described further below.

[0117] For the specified functions (e.g., DFLTCC-QAF, DFLTCC-GDHT, DFLTCC-CMPR, DFLTCC-XPND), the contents of general registers 0, 1 and R3 are not modified.

- 5 Further, in one example, R1 field 304 designates an even-odd pair of general registers. It is to designate an even-numbered register and is not to designate general register 0; otherwise, a specification exception is recognized.

- [0118] As depicted in FIGS. 3E-3F and described in further detail herein, the contents of general register R1 318 indicate a first operand address 320, and the contents of general  
 10 register R1 + 1 322 are used to determine the length 324 of the first operand. For instance, when the specified function is DFLTCC-CMPR or DFLTCC-XPND, the contents of general register R1 318 specify a logical address of the leftmost byte of the first operand. When the specified function is DFLTCC-CMPR, the contents of general register R1 + 1, in conjunction with the values of the new task (NT) and sub-byte boundary (SBB) fields of the parameter  
 15 block (described below), specify the length of the first operand. The following table provides examples which demonstrate the length of the first operand for the DFLTCC-CMPR function as a function of the contents of general register R1 + 1, the NT field, and the SBB field:

| <u>Contents of GR R1 + 1 (hex)</u> | <u>NT</u> | <u>SBB (binary)</u> | <u>Length of First<br/>Operand</u> |
|------------------------------------|-----------|---------------------|------------------------------------|
| 00000000 00000002                  | 0         | 001                 | 15 bits                            |
| 00000000 00000001                  | 1         | ---                 | 8 bits                             |
| 00000000 00000001                  | 0         | 000                 | 8 bits                             |
| 00000000 00000001                  | 0         | 011                 | 5 bits                             |
| 00000000 00000001                  | 0         | 111                 | 1 bit                              |
| 00000000 00000000                  | -         | ---                 | 0 bits                             |

[0119] When the specified function is DFLTCC-XPND, the contents of general register R1 + 1 specify the length of the first operand. When the specified function is DFLTCC-CMPR or DFLTCC-XPND, the results of compressing or decompressing data are stored at the first operand location. When the DFLTCC-QAF or DFLTCC-GDHT function is specified, the contents of general registers R1 and R1 + 1 are ignored.

[0120] Moreover, for the specified functions (e.g., DFLTCC-QAF, DFLTCC-GDHT, DFLTCC-CMPR and DFLTCC-XPND), in one example, R2 field 306 designates an even-odd pair of general registers. It is to designate an even-numbered register and is not to designate general register 0; otherwise, a specification exception is recognized.

[0121] As depicted in FIGS. 3G-3H and described in further detail herein, the contents of general register R2 326 indicate a second operand address 328, and the contents of general register R2 + 1 330 are used to determine the length 332 of the second operand. For instance, when the specified function is DFLTCC-GDHT, DFLTCC-CMPR, or DFLTCC-XPND, the contents of general register R2 specify a logical address of the leftmost byte of the second operand. When the specified function is DFLTCC-CMPR or DFLTCC-GDHT, the contents of general register R2 + 1 specify the length of the second operand. When the specified function is DFLTCC-XPND, the contents of general register R2 + 1, in conjunction with the values of the NT and SBB fields of the parameter block, specify the length of the second operand. When the second operand length is referenced and has a non-zero value at

the beginning of the execution of the instruction, data is fetched from the second operand location. When the second operand length is referenced, has a value of zero at the beginning of the execution of the instruction, and the continuation flag (CF) field of the parameter block is one at the beginning of the execution of the instruction, the second operand is not accessed.

- 5 [0122] When the DFLTCC-QAF function is specified, the contents of general registers R2 and R2 + 1 are ignored. When the DFLTCC-GDHT function is specified and the contents of general register R2 + 1 specify a length equal to zero, a specification exception is recognized, and the second operand is not accessed. When the DFLTCC-CMPR or DFLTCC-XPND function is specified, the continuation flag (CF) field of the parameter block
- 10 is zero at the beginning of the execution of the instruction, and the contents of general register R2 + 1 specify a length equal to zero, a specification exception is recognized, and the second operand is not accessed.

- [0123] As shown in FIG. 3I, when the specified function is DFLTCC-CMPR or DFLTCC-XPND and the history buffer type (HBT) is circular (e.g. HBT 310=1), the contents
- 15 of general register R3 335 specify a circular history buffer address 337. For instance, a logical address of the leftmost byte of the third operand is specified. It is to designate, e.g., a 4 K-byte boundary; otherwise a specification exception is recognized. In one example, the circular history buffer is located at the third operand location. When the specified function is DFLTCC-CMPR or DFLTCC-XPND and the HBT is zero, the contents of general register R3
- 20 are ignored. When the DFLTCC-QAF or DFLTCC-GDHT function is specified, the contents of general register R3 are ignored. For the specified functions (e.g., DFLTCC-QAF, DFLTCC-GDHT, DFLTCC-CMPR and DFLTCC-XPND), the R3 field is not to designate general register 0 or general register 1; otherwise, a specification exception is recognized, in one example.

- 25 [0124] As part of the operation, when the specified function is DFLTCC-CMPR, the address in general register R1 is incremented by the number of bytes processed of the first operand that included processing bit position 0, and the length in general register R1 + 1 is decremented by the same number; the address in general register R2 is incremented by the number of bytes processed of the second operand, and the length in general register R2 + 1 is
- 30 decremented by the same number. The number of bytes processed of the first operand that

included processing bit position 0 is, for instance, the integer quotient resulting from an integer division with the dividend being the sum of the number of output bits processed and the original value of the SBB, and the divisor being a value of eight. The formation and updating of the addresses and lengths are dependent on the addressing mode, as described below.

[0125] As part of the operation, when the specified function is DFLTCC-XPND, the address in general register R1 is incremented by the number of bytes processed of the first operand, and the length in general register R1 + 1 is decremented by the same number; the address in general register R2 is incremented by the number of bytes processed of the second operand that included processing bit position 0, and the length in general register R2 + 1 is decremented by the same number. The number of bytes processed of the second operand that included processing bit position 0 is the integer quotient resulting from an integer division with the dividend being the sum of the number of input bits processed and the original value of the SBB, and the divisor being a value of eight. The formation and updating of the addresses and lengths are dependent on the addressing mode, as described below.

[0126] In the 24-bit addressing mode, the following apply, in one embodiment:

[0127] • The contents of bit positions 40-63 of general registers 1, R1, R2, and R3 constitute the addresses of the parameter block, first operand, second operand, and circular history buffer, respectively, and the contents of bit positions 0-39 are ignored.

[0128] • Bits 40-63 of the updated first operand and second operand addresses replace the corresponding bits in general registers R1 and R2, respectively. Carries out of bit position 40 of the updated addresses are ignored, and the contents of bit positions 32-39 of general registers R1 and R2 are set to zeros. The contents of bit positions 0-31 of general registers R1 and R2 remain unchanged. When the instruction ends with partial or normal completion, and an updated operand address equals the operand address at the beginning of execution of the instruction, bit positions 32-39 of the corresponding general register are set to zeros.

[0129] • The contents of bit positions 32-63 of general registers  $R1 + 1$  and  $R2 + 1$  form, e.g., 32-bit unsigned binary integers which specify the number of bytes in the first and second operands, respectively. The contents of bit positions 0-31 of general registers  $R1 + 1$  and  $R2 + 1$  are ignored.

5 [0130] • Bits 32-63 of the updated first operand and second operand lengths replace the corresponding bits in general registers  $R1 + 1$  and  $R2 + 1$ , respectively. The contents of bit positions 0-31 of general registers  $R1 + 1$  and  $R2 + 1$  remain unchanged.

[0131] In the 31-bit addressing mode, the following apply, in one embodiment:

[0132] • The contents of bit positions 33-63 of general registers 1,  $R1$ ,  $R2$ , and  
10  $R3$  constitute the addresses of the parameter block, first operand, second operand, and circular history buffer, respectively, and the contents of bit positions 0-32 are ignored.

[0133] • Bits 33-63 of the updated first operand and second operand addresses replace the corresponding bits in general registers  $R1$  and  $R2$ , respectively. Carries out of bit position 33 of the updated addresses are ignored, and the content of bit position 32 of general  
15 registers  $R1$  and  $R2$  is set to zero. The contents of bit positions 0-31 of general registers  $R1$  and  $R2$  remain unchanged. When the instruction ends with partial or normal completion, and an updated operand address equals the operand address at the beginning of the execution of the instruction, bit position 32 of the corresponding general register is set to zero.

[0134] • The contents of bit positions 32-63 of general registers  $R1 + 1$  and  $R2 + 1$  form 32-bit unsigned binary integers which specify the number of bytes in the first and  
20 second operands, respectively. The contents of bit positions 0-31 of general registers  $R1 + 1$  and  $R2 + 1$  are ignored.

[0135] • Bits 32-63 of the updated first operand and second operand lengths replace the corresponding bits in general registers  $R1 + 1$  and  $R2 + 1$ , respectively. The  
25 contents of bit positions 0-31 of general registers  $R1 + 1$  and  $R2 + 1$  remain unchanged.

[0136] In the 64-bit addressing mode, the following apply, in one embodiment:

[0137] • The contents of bit positions 0-63 of general registers 1, R1, R2, and R3 constitute the addresses of the parameter block, first operand, second operand, and circular history buffer, respectively.

[0138] • Bits 0-63 of the updated first operand and second operand addresses replace the corresponding bits in general registers R1 and R2, respectively. Carries out of bit position 0 of the updated addresses are ignored.

[0139] • The contents of bit positions 0-63 of general registers R1 + 1 and R2 + 1 form 64-bit unsigned binary integers which specify the number of bytes in the first and second operands, respectively.

10 [0140] • Bits 0-63 of the updated first operand and second operand lengths replace the corresponding bits in general registers R1 + 1 and R2 + 1, respectively.

[0141] In the access-register mode, access registers 1, R1, R2, and R3 specify the address spaces containing the parameter block, first operand, second operand, and circular history buffer, respectively. When DFTCC-CMPR with an in-line history buffer is specified  
15 in the access-register mode, access register R2 specifies the address space containing the in-line history. When DFTCC-XPND with an in-line history buffer is specified in the access-register mode, access register R1 specifies the address space containing the in-line history.

[0142] Further details regarding the various functions are described below:

[0143] Function Code 0: DFLTCC-QAF (Query Available Functions)

20 [0144] The DFLTCC-QAF (query available functions) function provides a mechanism to indicate the availability of the installed functions and installed parameter block formats. One example format of a parameter block for the DFLTCC-QAF function is described with reference to FIG. 3J. In one example, a parameter block 340 for the DFLTCC-QAF function (e.g., function code 0) includes an installed functions vector 342 and an  
25 installed parameter block formats vector 346. In one particular example, these vectors are stored to bytes 0-15 and bytes 24-25, respectively, of the parameter block. Each of the vectors is further described below.

[0145] As an example, bits 0-127 of installed functions vector 342 correspond to function codes 0-127, respectively, of the DEFLATE Conversion Call instruction. When a bit is, e.g., one, the corresponding function is installed; otherwise, the function is not installed.

[0146] Further, in one example, bits 0-15 of installed parameter block formats vector 346 correspond to parameter block formats 0-15, respectively for the DFLTCC-GDHT, DFLTCC-CMPR and DFLTCC-XPND functions. When a bit is, e.g., one, the corresponding parameter block format is installed; otherwise, the parameter block format is not installed. In one example, zeros are stored to reserved bytes 16-23 and 26-31 of the parameter block.

[0147] Although certain fields are described with respect to parameter block 340, additional, fewer and/or other fields may be included in other embodiments.

[0148] The contents of general registers R1, R2, R3, R1 + 1, and R2 + 1 are ignored by the DFLTCC-QAF function, in one embodiment.

[0149] A PER (program event recording) storage alteration event is recognized, when applicable, for the parameter block. A PER zero address detection event is recognized, when applicable, for the parameter block.

[0150] In one example, condition code 0 is set when execution of the DFLTCC-QAF function completes; condition codes 1, 2, and 3 are not applicable to the query function, in one example.

[0151] Function Code 1: DFLTCC-GDHT (Generate Dynamic-Huffman Table)

[0152] When the DFLTCC-GDHT function is specified, the second operand is used, for instance, as a source to generate a compressed representation of a dynamic-Huffman table (DHT), as specified by the DEFLATE standard.

[0153] In one example, the DFLTCC-GDHT function uses a parameter block, an example of which is described with reference to FIG. 3K. In the example parameter block described herein, specific locations within the parameter block for specific fields and specific sizes of the fields are indicated (e.g., specific bytes and/or bits). However, other locations and/or sizes may be provided for one or more of the fields. Further, although the setting of a

bit to a particular value e.g., one or zero, is specified, this is only an example. The bit may be set to a different value, such as the opposite value or to another value, in other examples. Many variations are possible.

[0154] Additionally, in one example, the parameter block includes one or more preserved fields and one or more reserved fields. Preserved fields are not modified by the DFLTCC-GDHT function. Preserved fields are distinguished from reserved fields to enable a program to initialize a single storage location, use that storage location for the parameter block of a DFLTCC-GDHT function, and subsequently use the same storage location for the parameter block of a DFLTCC-CMPR function. Reserved fields are to contain zeros; otherwise, the program may not operate compatibly in the future. When an operation ends, reserved fields may be stored as zeros or may remain unchanged.

[0155] Yet further, some of the fields are used by other functions (e.g., DFLTCC-CMPR or DFLTCC-XPND), and thus, aspects relating to those functions may also be described with the description of those fields.

[0156] In one example, a parameter block 360 for the DFLTCC-GDHT function includes the following fields:

[0157] Parameter Block Version Number (PBVN) 362: Bytes 0-1 of the parameter block specify the version and size of the parameter block. Bits 0-11 of the PBVN are reserved and should contain zeros; otherwise, the program may not operate compatibly in the future. Bits 12-15 of the PBVN contain an unsigned binary integer specifying the format of the parameter block. The DFLTCC-QAF function provides a mechanism to indicate the parameter block formats available. When the format of the parameter block specified is not supported by the model, a general operand data exception is recognized. The PBVN is specified by the program and is not modified during execution of the instruction.

[0158] Model Version Number (MVN) 363: Byte 2 of the parameter block is an unsigned binary integer identifying the model which executed the instruction. The program is not required to initialize the MVN. The MVN is updated during execution of the instruction. The value stored in the MVN is model-dependent.

[0159] Dynamic-Huffman Table (DHT) Generation Control (DHTGC) 364: Bit 2 of byte 17 of the parameter block applies to generating a dynamic-Huffman table (DHT). The DHT specifies Huffman codes for symbols representing literal bytes, duplicate string lengths, end-of-block (EOB) symbol, and duplicate string pointer distances. The value of a Huffman code for a particular symbol is a function of the count of occurrences for the entity, which the symbol represents, in the uncompressed form of the data. When the count for a symbol is zero, there is no Huffman code in the DHT for the symbol. The DHTGC specifies counts equal to zero will be treated as follows, in one example:

|  | DHTGC | Meaning |
|--|-------|---------|
|--|-------|---------|

|    |   |                                                                                                                                          |
|----|---|------------------------------------------------------------------------------------------------------------------------------------------|
| 10 | 0 | Treat counts of literal bytes, duplicate string lengths, and pointer distances equal to zero as equal to one (generate a universal DHT). |
|----|---|------------------------------------------------------------------------------------------------------------------------------------------|

|  |   |                                                                                               |
|--|---|-----------------------------------------------------------------------------------------------|
|  | 1 | Treat counts of duplicate string lengths and pointer distances equal to zero as equal to one. |
|--|---|-----------------------------------------------------------------------------------------------|

15 [0161] A DHT which specifies a Huffman code for every possible value of literal bytes, an EOB symbol, duplicate string lengths, and duplicate string pointer distances is called a universal DHT. A DHT which does not specify Huffman codes for values of literal bytes, duplicate string lengths, or duplicate string pointer distances which do not occur in the uncompressed form of the data is called a non-universal DHT.

20 [0162] For all values of the DHTGC, the resulting DHT specifies Huffman codes for all possible duplicate string lengths and pointer distances, as defined by the DEFLATE standard. Therefore, the HLIT (Huffman literal) and HDIST (Huffman distances) sub-elements of the resulting compressed form of the DHT, described further below, each contain the value of, for instance, 29.

25 [0163] The DHTGC is an input to the operation when the DFLTCC-GDHT function is specified. The DHTGC does not apply to the operation when the DFLTCC-CMPR or DFLTCC-XPND function is specified. The DHTGC is not modified, in one embodiment, during execution of the instruction.

[0164] Operation Ending Supplemental Code (OESC) 365: Byte 19 of the parameter block is an unsigned binary integer providing additional information on the condition being reported to the program. Since this field is used by multiple functions, some of the conditions refer to fields of a parameter block used by other functions (e.g., the parameter block of FIG. 3L used by the DFLTCC-CMPR and DFLTCC-XPND functions). When the condition being reported is a general operand data exception, the operation is considered suppressed, although the OESC field of the parameter block is updated, in which case it is defined as follows, in one example:

| [0165] | <u>OESC (hex)</u> | <u>Meaning</u>                                                                                                                                                        |
|--------|-------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 10     | 00                | No additional information is provided.                                                                                                                                |
|        | 01                | The format of the parameter block, as specified by parameter block version number 362, is not supported by the model.                                                 |
| 15     | 02                | The DFLTCC-CMPR or DFLTCC-XPND function is specified, a history length field 385 (FIG. 3L) is greater than, e.g., 32,768, and a new task field 374 (FIG. 3L) is zero. |
|        | 11                | A compressed data block with BTYPE (Block Type) equal to 11 binary is encountered.                                                                                    |
|        | 12                | A compressed data block with BTYPE equal to 00 binary and NLEN not equal to the one's complement of LEN (length) is encountered.                                      |
| 20     | 21                | CDHTL field 366 (FIG. 3L) applies and is, e.g., less than 42 or greater than 2283.                                                                                    |
|        | 22                | The HLIT sub-element of a compressed DHT used during the operation is, e.g., greater than 29 (invalid DHT).                                                           |
| 25     | 23                | The HDIST sub-element of a compressed DHT used during the operation is, e.g., greater than 29 (invalid DHT).                                                          |

24 A compressed DHT used during the operation specifies a code which is in the sequence of codes specifying the bit lengths for the possible code lengths (e.g., 19) defined for a compressed DHT and is less than the length required by the Huffman algorithm to specify a functional Huffman tree (invalid DHT).

5

26 A compressed DHT used during the operation specifies code length 16 (copy previous code length) as the first code length for the set of elements consisting of literal bytes, an EOB symbol, and duplicate string lengths (invalid DHT).

27 A compressed DHT used during the operation specifies a code which is in the sequence of codes specifying code lengths for literal bytes, and the code does not match any of the codes determined to represent the set of referenced code lengths, as specified earlier in the compressed DHT (invalid DHT).

10

28 A compressed DHT used during the operation specifies a code which assigns code length 0 (CL0) to the EOB symbol. In this case, the corresponding DHT does not specify a Huffman code to represent an EOB symbol (invalid DHT).

15

29 A compressed DHT used during the operation specifies a code which is in the sequence of codes specifying code lengths for duplicate string lengths and pointer distances, and the code does not match any of the codes determined to represent the set of referenced code lengths, as specified earlier in the compressed DHT (invalid DHT).

20

2A A compressed DHT used during the operation specifies a number of code lengths which is greater than the number of Huffman codes in the DHT, as specified by the sum of the values in the HLIT field, the HDIST field, and, e.g., 258. This is possible with an improper uses of code lengths 16, 17, and 18, as examples (invalid DHT).

25

2B                      A compressed DHT used during the operation specifies a code length for the set of literal bytes, EOB symbol, and duplicate string lengths, which is less than the length required by the Huffman algorithm to specify a functional Huffman tree (invalid DHT).

5                      2D                      A compressed DHT used during the operation specifies a code length for the set of duplicate string pointer distances, which is less than the length required by the Huffman algorithm to specify a functional Huffman tree (invalid DHT).

10                      2F                      CDHTL field 366 (FIG. 3L) applies and does not equal the length of the compressed DHT in CDHT field 367 (FIG. 3L) used during the operation.

15                      31                      A compressed DHT used during the operation does not specify a Huffman code corresponding to a literal byte or a duplicate string length processed during the operation (deficient non-universal DHT), or the DFLTCC-XPND function is specified and a compressed data symbol, which is encountered in a compressed data block with BTYPE equal 01 binary, specifies an invalid code for a duplicate string length (11000110 or 11000111 binary).

20                      32                      A compressed DHT used during the operation does not specify a Huffman code corresponding to a duplicate string pointer distance processed during the operation (deficient non-universal DHT), or the DFLTCC-XPND function is specified and a compressed data symbol, which is encountered in a compressed data block with BTYPE equal 01 binary, specifies an invalid code for a duplicate string pointer distance (11110 or 11111 binary).

25                      40                      A compressed data symbol is encountered which is a duplicate string pointer and specifies a distance greater than the length of history available at the point of processing the symbol.

[0166]                      When the operation ends without reporting a general operand data exception, zeros are stored to the OESC field.

[0167] Support for supplemental codes other than zero is model-dependent. When multiple conditions exist, it is model-dependent which code, if any, is reported in the OESC field.

[0168] Compressed Dynamic-Huffman Table Length (CDHTL) 366: Twelve bits, starting with bit 4 of byte 56 through bit 7 of byte 57, of the parameter block contain an unsigned binary integer which specifies the length, as a bit count, of the compressed format of the DHT in the CDHT field of the parameter block (e.g., CDHT 367).

[0169] The CDHTL is an output from the operation when the DFLTCC-GDHT function is specified.

10 [0170] The CDHTL is an input to the operation when the DFLTCC-CMPR function is specified and the Huffman table type (e.g., HTT 376 of FIG. 3L) is one. When the CDHTL does not specify an appropriate length for the CDHT, a general operand data exception is recognized. The CDHTL is not modified when the DFLTCC-CMPR function is specified.

[0171] When the DFLTCC-XPND function is specified and the operation ends after decoding only a portion of a block with BTYPE 10 binary, the length of the compressed representation of the DHT in the block is stored to this field. When the DFLTCC-XPND function is specified and the operation ends at a block boundary or after decoding only a portion of a block with BTYPE 00 or 01 binary, zeros are stored to this field. When an uncompressing operation is resumed within a block with BTYPE 10 binary (that is when CF (continuation flag 373 of FIG. 3L) equals one and IFS (incomplete function status 383) equals C or D hex, described below), this field is an input to the operation.

[0172] Compressed Dynamic-Huffman Table (CDHT) 367: Bytes 64-351 of the parameter block contain a compressed format of a dynamic-Huffman table (DHT).

[0173] The DHT specifies Huffman codes (bit sequences) to represent two sets of elements. The elements for one set include literal bytes, an EOB symbol, and duplicate string lengths. The elements for the other set include duplicate string pointer distances. The compressed representation of the DHT defines a set of code lengths and specifies a code length (CL) for each element of each set. The Huffman code for an element expected to be

referenced during an operation is derived from the CL specified for that element and the number of elements in the same set with the same specified CL. Specifically, the compressed representation of the DHT includes the following, as an example:

[0174] • An HLIT field to specify the number of Huffman codes representing  
5 literal bytes, an EOB symbol and duplicate string lengths.

[0175] • An HDIST field to specify the number of Huffman codes representing  
duplicate string pointer distances.

[0176] • An HCLLEN (Huffman code lengths) field to specify the number of  
Huffman codes representing code lengths.

10 [0177] • A sequence of codes specifying a bit length for each of the, for  
instance, 19 code lengths defined for the compressed DHT.

[0178] • A sequence of codes specifying a code length for each of the elements  
of the set consisting of literal bytes, an EOB symbol, and duplicate string lengths.

[0179] • A sequence of codes specifying a code length for each of the elements  
15 of the set consisting of duplicate string pointer distances.

[0180] Further details of a compressed representation of a DHT are described below  
with reference to the description of a compressed data block with block type 10 binary.

[0181] In one example, the compressed representation of the DHT is left justified in  
the CDHT field. That is, the rightmost bit of byte 64 contains the least-significant bit of the  
20 HLIT sub-element of the compressed representation of the DHT.

[0182] The compressed representation of a DHT is an output from the operation when  
the DFLTCC-GDHT function is specified.

[0183] The compressed representation of a DHT is an input to the operation when the  
DFLTCC-CMPR function is specified and HTT, described below, is one. The CDHT field is  
25 not modified by the DFLTCC-CMPR function.

[0184] When the DFLTCC-XPND function is specified and the operation ends after decoding only a portion of a block with BTYPE 10 binary, the compressed representation of the DHT in the block is stored to this field. When the DFLTCC-XPND function is specified and the operation ends at a block boundary or after decoding only a portion of a block with BTYPE 00 or 01 binary, zeros are stored to this field. When an uncompressing operation is resumed within a block with BTYPE 10 binary (that is when CF equals one and IFS equals C or D hex), this field is an input to the operation.

[0185] When the CDHT is modified, bits of the field not used to represent the compressed representation of the DHT are stored as zeros.

10 [0186] Although various fields are described above with respect to parameter block 360, additional, fewer and/or other fields may be included in other embodiments.

[0187] Aspects of DHT generation are specified by the program to the machine using dynamic-Huffman table generation control (DHTGC) field 364 of the parameter block. It is intended that the source contains uncompressed data and subsequent to completing the operation, the generated result is specified with the DFLTCC-CMPR function to compress the same source.

[0188] In one embodiment, there is no history to reference from prior operations while processing the current operation.

[0189] When the contents of general register R2 + 1 specify a length greater than, e.g., 32 K-bytes, the following applies, in one example:

[0190] • Only the first 32 K-bytes of the second operand are used to generate the DHT.

[0191] • Access exceptions are not recognized for locations beyond the first 32 K-bytes of the second operand.

25 [0192] When the contents of general register R2 + 1 specify a length equal to zero, a specification exception is recognized, and the second operand is not accessed.

[0193] The resulting compressed DHT includes a Huffman code representing an end-of-block (EOB) symbol.

[0194] A compressed format of the generated DHT is stored to compressed dynamic-Huffman Table (CDHT) field 367 of the parameter block. The length of the compressed  
5 format of the generated DHT is stored to CDHTL field 366 of the parameter block.

[0195] The operation includes storing a model identification to a model version number field 363 of the parameter block.

[0196] When the operation ends without recognizing a general operand data exception, zeros are stored to operation ending supplemental code (OESC) field 365 of the  
10 parameter block.

[0197] Condition code 0 is set when execution of the DFLTCC-GDHT function completes; condition codes 1, 2 and 3 are not applicable to the DFLTCC-GDHT function.

[0198] General registers R2 and R2 + 1 are not modified by the operation.

[0199] The contents of general registers R1, R1 + 1 and R3 are ignored when the  
15 DFLTCC-GDHT function is specified.

[0200] A PER zero address detection event is recognized, when applicable, for the second operand location and for the parameter block.

[0201] Function Code 2: DFLTCC-CMPR (Compress)

[0202] When the DFLTCC-CMPR function is specified, a compressing operation is  
20 performed. The operation includes encoding data from the second operand location into compressed data symbols, which are stored to the first operand location.

[0203] In one example, the DFLTCC-CMPR function uses a parameter block, an example of which is described with reference to FIG. 3L. Some of the fields have been described above with respect to parameter block 360, and therefore, are listed below with the  
25 same reference number and not further described in detail.

[0204] In one example, parameter block 370 includes:

[0205] Parameter Block Version Number (PBVN) 362.

[0206] Model Version Number (MVN) 363.

[0207] Continuation Flag (CF) 373: Bit 63 of the parameter block, when one,  
5 indicates the operation is partially complete and the contents of the continuation state buffer (e.g., in continuation state buffer field 392) may be used to resume the operation. The program is to initialize the continuation flag (CF) to zero and not modify CF in the event the instruction is to be re-executed for the purpose of resuming the operation; otherwise results are unpredictable.

10 [0208] New Task (NT) 374: Bit 0 of byte 16 of the parameter block, when one, indicates the operation applies to the beginning of a compressed data set. Therefore, no history and no check value from a prior operation applies to the current operation. When NT is one at the beginning of the operation, and the operation ends after partial completion, zero is stored to the NT field. When NT is zero, history and a check value from a prior operation  
15 apply to the current operation.

[0209] Check Value Type (CVT) 375: Bit 2 of byte 16 of the parameter block specifies the type of check value contained in the check value field of the parameter block (e.g., field 387). When CVT is zero, the check value type is, e.g., a 32-bit cyclic-redundancy-check (CRC-32). When CVT is one, the check value type is, e.g., a 32-bit Adler checksum  
20 (Adler-32). The CVT bit is not modified during execution of the instruction.

[0210] Huffman Table Type (HTT) 376: Bit 4 of byte 16 of the parameter block, when zero, specifies a table containing fixed-Huffman codes (FHT), as defined by the DEFLATE standard, is used during a compression operation. When the HTT is one, a table containing dynamic-Huffman codes (DHT), as specified in the CDHT field of the parameter  
25 block, is used during a compression operation. The HTT does not apply to decompressing operations. The HTT bit is not modified during execution of the instruction.

[0211] Block Continuation Flag (BCF) 377: Bit 5 of byte 16 of the parameter block applies when the DFLTCC-CMPR function is specified. When zero, a 3-bit block header, and when applicable, the compressed format of a dynamic-Huffman table, as specified in the CDHT field of the parameter block (e.g., field 367), is stored to the first operand location  
5 prior to storing any compressed data elements. When one, neither a block header nor a compressed format of a DHT is stored to the first operand location. When NT is one, BCF is treated as equal to zero. The BCF bit is not modified during execution of the instruction.

[0212] Block Closing Control (BCC) 378: Bit 6 of byte 16 of the parameter block applies when the DFLTCC-CMPR function is specified. When one, subsequent to storing all  
10 compressed data symbols, an end-of block (EOB) symbol is stored to the first operand location. When the HTT specifies using an FHT, Huffman code 0000000 binary (which corresponds to the intermediate integer representation of 256 in the table specifying codes for literal bytes, an EOB symbol, and duplicate string lengths), as an example, is used for the EOB symbol. When the HTT specifies using a DHT, the Huffman code for the EOB symbol  
15 is specified in the DHT. When the BCC bit is zero, an EOB symbol is not stored to the first operand location. The BCC bit is not modified during execution of the instruction.

[0213] Block Header Final (BHF) 379: Bit 7 of byte 16 of the parameter block applies when the DFLTCC-CMPR function is specified and either BCF 377 is zero or NT 374 is one; otherwise the BHF does not apply. When applicable and one, the first bit of the block  
20 header (BFINAL) is set to one before storing the block header to the first operand location. When applicable and zero, the first bit of the block header (BFINAL) is set to zero before storing the block header to the first operand location. The BHF bit is not modified during execution of the instruction.

[0214] DHT Generation Control (DHTGC) 364: The DHTGC does not apply to the  
25 operation when the DFLTCC-CMPR function is specified. The DHTGC is not modified during execution of the instruction.

[0215] Sub-Byte Boundary (SBB) 381: Bits 5-7 of byte 18 of the parameter block contain an unsigned binary integer specifying the boundary between processed and unprocessed bits within a byte of the compressed data stream. The byte of the stream

referenced is the last byte referenced, meaning the rightmost byte, when an operation ends, and is the first byte to be referenced, meaning the leftmost byte, when an operation begins or resumes. When the DFLTCC-CMPR function is specified, the SBB applies to the byte designated by the first operand address. When the DFLTCC-XPND function is specified, the SBB applies to the byte designated by the second operand address. The SBB specifies the number of rightmost bits that have been processed. The SBB is an input to the operation and an output of the operation.

[0216] One example of a compressed data stream when SBB has a value of 011 binary is depicted in FIG. 4. The data that has been processed after the end of the operation is depicted at 400; and the data to be processed before the start of the operation is depicted at 402.

[0217] Further, FIGS. 5A-5C provide examples which demonstrate how the SBB applies to the DFLTCC-CMPR function. For instance, one example of how SBB applies before and after executing the DFLTCC-CMPR function is depicted in FIG. 5A. Other examples are depicted in FIGS. 5B-5C. When NT 374 is one, SBB 381 is treated as equal to 000 binary.

[0218] Returning to FIG. 3L, additional fields of parameter block 370 are described:

[0219] Operation Ending Supplemental Code (OESC) 365.

[0220] Incomplete Function Status (IFS) 383: Bits 4-7 of byte 21 of the parameter block contain status information when certain operations end. When a decompressing operation ends, the IFS conveys information about the second operand as follows, in on example:

| [0221] | <u>IFS (binary)</u> | <u>Meaning</u>                                                                           |
|--------|---------------------|------------------------------------------------------------------------------------------|
|        | 0000                | The operation ended after decoding the last element of a block with BFINAL equal to one. |

1000 The operation ended after decoding an element, other than the last element, of a block with BTYPE equal 00 binary and BFINAL equal to zero.

5 1001 The operation ended after decoding an element, other than the last element, of a block with BTYPE equal 00 binary and BFINAL equal to one.

1010 The operation ended after decoding an element, other than the last element, of a block with BTYPE equal 01 binary and BFINAL equal to zero.

10 1011 The operation ended after decoding an element, other than the last element, of a block with BTYPE equal 01 binary and BFINAL equal to one.

15 1100 The operation ended after decoding an element, other than the last element, of a block with BTYPE equal 10 binary and BFINAL equal to zero.

1101 The operation ended after decoding an element, other than the last element, of a block with BTYPE equal 10 binary and BFINAL equal to one.

20 1110 The operation ended at a block boundary, the last element of a block with BFINAL equal to one has not been decoded, and the block header first element of the subsequent block has not yet been processed.

[0222] In one embodiment, an uncompressing operation may end with IFS equal 0000 binary and not satisfy normal completion. In such cases, the operation ends with condition code 1 or 3 set.

25 [0223] When a compressing operation ends, the IFS field is undefined, but may be modified.

[0224] The IFS is not an input to the operation.

[0225] Incomplete Function Length (IFL) 384: Bytes 22-23 of the parameter block contain length information when certain operations end. For a decompressing operation, the IFL applies to the second operand. When a decompressing operation ends after decoding  
5 some, but not all of a block with BTYPE equal 00 binary, the IFL contains an unsigned binary integer specifying the number of bytes of the block in the second operand, which have not yet been processed. Bytes 22-23 contain the IFL in, for instance, big-endian byte order unlike the LEN field of a block with BTYPE equal 00 binary, which is in, for instance, little-endian byte order.

10 [0226] When a decompressing operation ends after decoding a complete block with BTYPE equal 00 binary and BFINAL equal to one, zeros are stored to the IFL field. When a decompressing operation ends after decoding some, but not all of a block with a non-zero BTYPE, or ends at a block boundary, the IFL field is undefined, but may be modified.

[0227] When a compressing operation ends, the IFL field is undefined, but may be  
15 modified.

[0228] The IFL is not an input to the operation.

[0229] History Length (HL) 385: Bytes 44-45 of the parameter block contain an unsigned binary integer specifying the number of bytes of history in the history buffer which can be referenced during an operation. The HL applies to in-line and circular history buffers.  
20 When new task (NT) equals one, no history applies to the beginning of the operation and the history length is treated as zero as an input to the operation.

[0230] A general operand data exception is recognized when the history length is greater than, e.g., 32,768 and NT equals zero.

[0231] The history length is modified during compressing and decompressing  
25 operations. When the sum of the original HL and the number of uncompressed data bytes processed during the operation is less than, or equal to, e.g., 32,768, the updated HL is equal

to the sum of the original HL and the number of uncompressed data bytes processed during the operation; otherwise the updated HL is equal to the value of 32,768.

[0232] History Offset (HO) 386: Fifteen bits, starting with bit 1 of byte 46, through bit 7 of byte 47, of the parameter block, contain an unsigned binary integer specifying an offset in the third operand when the history buffer type is circular. The sum of the contents of R3 and the history offset designates the location of the first byte of history within the circular history buffer, which is the least recently processed byte of uncompressed data in the buffer. When the history buffer type is circular, history offset is an input to the operation and is updated at the end of the operation. When the sum of the original HL and the number of uncompressed data bytes processed during the operation is less than, or equal to, e.g., 32,768, the updated HO is equal to the original HO; otherwise, the updated HO is equal to the sum of the original HO, the original HL, and the number of uncompressed data bytes processed during the operation, modulo 32,768.

[0233] When the history buffer type is in-line, the HO field of the parameter block is undefined, but may be modified.

[0234] Check Value 387: Bytes 48-51 of the parameter block contain a check value. As part of the operation, a check value is generated. The check value applies to the uncompressed data operand. That is, the check value applies to the second operand for the DFLTCC-CMPR function and applies to the first operand for the DFLTCC-XPND function. When CVT bit 375 is zero, a, e.g., 32-bit cyclic-redundancy-check check value (CRC-32) is generated. When the CVT bit is one, a, e.g., 32-bit Adler checksum check value (Adler-32) is generated.

[0235] The inputs to generating a check value are, for instance, a 4-byte base and the uncompressed data processed during the operation. The base input provides the means to compute a single and consistent check value for a set of compressed data blocks, regardless of the number of times the DFLTCC instruction is executed to process the complete set of compressed data blocks. When the NT bit is zero, the original value in the check value field is used for the base input in generating a check value.

[0236] When an Adler-32 check value is generated, the following apply, in one example:

[0237] • When the NT bit is one, a value of one is used for the 4-byte base input.

5 [0238] • The sums defined in the Adler-32 check value generation are modulo 65,521.

[0239] • The result is stored to the check value field in big-endian byte order. That is, the most significant byte of the check value is located in byte 48 and the least significant byte of the check value is located in byte 51.

10 [0240] When a CRC-32 check value is generated, the following apply, in one embodiment:

[0241] • When the NT bit is one, a value of zero is used for the 4-byte base input.

15 [0242] • The polynomial used as the divisor in generating a CRC-32 check value is  $x^{32} + x^{26} + x^{23} + x^{22} + x^{16} + x^{12} + x^{11} + x^{10} + x^8 + x^7 + x^5 + x^4 + x^2 + x^1 + x^0$ , which is represented as 104C11DB7 hex. In this representation, the leftmost bit corresponds to the most significant bit.

[0243] • The first and final stages of generating the check value are computing the one's complement of the base input and computing the one's complement of the result,  
20 prior to storing the result, respectively.

[0244] • The result is stored to the check value field in little-endian byte order. That is, the least significant byte of the check value is located in byte 48 and the most significant byte of the check value is located in byte 51.

25 [0245] In one example, the check value is only meaningful to the program when the operation ends with condition code 0 set; otherwise, the check value is only an intermediate result and only meaningful to resume the operation. When the DFLTCC-CMPR function is

specified and the operation ends with condition code 1, 2, or 3 set some bytes to the left of the byte designated by the second operand address may not be included in the computation of the resulting check value. When the DFLTCC-XPND function is specified and the operation ends with condition code 1, 2, or 3 set, some result bytes not yet stored to the right of the byte  
5 designated by the first operand address may already be included in the computation of the resulting check value.

[0246] End-Of-Block Symbol (EOBS) 388: Fifteen bits, starting with bit 0 of byte 52, through bit 6 of byte 53, of the parameter block, contain an end-of-block (EOB) symbol. The end-of-block length (EOBL) field 389 of the parameter block specifies the length of the EOB  
10 symbol in the EOBS field. The EOB symbol is left justified in the EOBS field. Bits of the EOBS field not occupied by the EOB symbol are stored as zeros. The EOBS field is an output of the operation when compressing data, regardless of which type of Huffman table applies. The EOBS field is not used as an input to the operation.

[0247] Bit 0 of byte 52 contains the most significant bit of the EOB symbol. When the  
15 length of the EOB symbol is 7 bits, bit 6 of byte 52 contains the least significant bit of the EOB symbol. When the length of the EOB symbol is 15 bits, bit 6 of byte 53 contains the least significant bit of the EOB symbol.

[0248] For blocks using a FHT, the EOB symbol is 0000000 binary, as defined by the DEFLATE standard. For blocks using a DHT, the EOB symbol is defined by the DHT. The  
20 EOB symbol is conveyed in order to provide the capability for the program to close a block.

[0249] The EOBS field is undefined when the DFLTCC-XPND function is specified but may be modified.

[0250] End-Of-Block Length (EOBL) 389: Bits 0-3 of byte 54 of the parameter block contain an unsigned binary integer specifying the length of the end-of-block (EOB) symbol in  
25 EOBS field 388 of the parameter block. The length specifies the number of bits which the EOB symbol occupies in the EOBS field. The EOBL field is an output of the operation when compressing data, regardless of which type of Huffman table applies. The EOBL field is not used as an input to the operation.

[0251] The EOBL field is undefined when the DFLTCC-XPND function is specified, but may be modified.

[0252] Compressed Dynamic-Huffman Table Length (CDHTL) 366.

[0253] Compressed Dynamic-Huffman Table (CDHT) 367: The compressed  
5 representation of a DHT is an input to the operation when the DFLTCC-CMPR function is specified and HTT is one. The CDHT field is not modified by the DFLTCC-CMPR function.

[0254] Continuation State Buffer (CSB) 392: When conditions cause a value of one to be stored in CF field 373, internal state data is stored to bytes 384-1535 of the parameter block; otherwise, bytes 384-1535 of the parameter block are undefined and may be modified.  
10 The internal state data stored is model-dependent and may be used subsequently to resume the operation. It is expected, but not required, for the program to initialize the continuation state buffer to contain, e.g., all zeros. Subsequent to the instruction ending with a non-zero condition code set, and prior to re-executing the instruction for the purpose of resuming the operation, the program should not modify the continuation state buffer; otherwise results are  
15 unpredictable.

[0255] Although various fields are described above with respect to parameter block 370, additional, fewer and/or other fields may be included in other embodiments.

[0256] One example of the compress operation is described below with respect to Compressing Data.

20 [0257] Normal completion for the DFLTCC-CCMPR function occurs when the entire second operand is compressed and stored to the first operand location. When the operation ends due to normal completion, the following occurs, in one example:

[0258] • A model-dependent value is stored to model version number (MVN) field 363 of the parameter block.

25 [0259] • Continuation flag (CF) field 373 of the parameter block is set to zero.

[0260] • Sub-byte boundary (SBB) field 381 of the parameter block is updated.

- [0261] • End-of-block length (EOBL) 389 and end-of block symbol (EOBS) 388 fields of the parameter block are updated.
- [0262] • History length (HL) field 385 of the parameter block is updated.
- [0263] • The history offset (HO) field 386 of the parameter block is updated,  
5 when applicable.
- [0264] • Operation ending supplemental code (OESC) field 365 of the parameter block is set to zeros.
- [0265] • Check value field 387 of the parameter block is updated.
- [0266] • The address in general register  $R_1$  is incremented by the number of  
10 bytes processed of the first operand that included processing bit 0, and the length in general register  $R_1 + 1$  is decremented by the same number. The number of bytes processed of the first operand that included processing bit 0 is the integer quotient resulting from an integer division with the dividend being the sum of the number of output bits processed and the original value of the SBB, and the divisor being a value of eight.
- 15 [0267] • The address in general register  $R_2$  is incremented by the number of source bytes processed, and the length in general register  $R_2 + 1$  is decremented by the same number.
- [0268] • Condition code 0 is set.
- [0269] The formation and updating of the addresses and lengths are dependent on the  
20 addressing mode.
- [0270] When normal completion occurs, CSB field 392 of the parameter block is undefined after the operation ends.
- [0271] When a CPU-determined number of bytes have been processed, the operation ends and the following occurs, in one example:
- 25 [0272] • Continuation flag (CF) bit 373 in the parameter block is set to one.

- [0273] • Continuation state buffer (CSB) field 392 in the parameter block is updated.
- [0274] • Sub-byte boundary (SBB) field 381 of the parameter block is updated.
- [0275] • History length (HL) field 385 of the parameter block is updated.
- 5 [0276] • History offset (HO) field 386 of the parameter block is updated, when applicable.
- [0277] • Check value field 387 of the parameter block is updated.
- [0278] • A model-dependent value is stored to model version number (MVN) field 363 of the parameter block.
- 10 [0279] • End-of-block length (EOBL) 389 and end-of block symbol (EOBS) 388 fields of the parameter block are updated.
- [0280] • Operation ending supplemental code (OESC) field 365 of the parameter block is set to zeros.
- [0281] • The address in general register  $R_1$  is incremented by the number of  
15 bytes processed of the first operand that included processing bit 0, and the length in general register  $R_1 + 1$  is decremented by the same number. The number of bytes processed of the first operand that included processing bit 0 is the integer quotient resulting from an integer division with the dividend being the sum of the number of output bits processed and the original value of the SBB, and the divisor being a value of eight.
- 20 [0282] • The address in general register  $R_2$  is incremented by the number of source bytes processed, and the length in general register  $R_2 + 1$  is decremented by the same number.
- [0283] • Condition code 3 is set.
- [0284] The formation and updating of the addresses and lengths are dependent on the  
25 addressing mode.

[0285] The CPU-determined number of bytes depends on the model, and may be a different number each time the instruction is executed.

[0286] Subsequent to the instruction ending with condition code 3 set, it is expected the program does not modify any input or output specification for the instruction and  
5 branches back to re-execute the instruction to resume the operation.

[0287] In certain situations, despite ending the instruction with condition code 3 set, the parameter block and general registers are not updated. These situations may occur when the CPU performs a quiescing operation or CPU retry while executing the DEFLATE Conversion Call instruction. In these cases, the CPU-determined number of bytes processed is  
10 zero, data may have been stored to the first operand location, data may have been stored to the third operand location, when applicable, and corresponding change bits have been set.

[0288] The first operand length is insufficient to complete the operation when any of the following conditions apply, in one example:

[0289] • The first operand length, as specified by the contents of general register  
15  $R_1 + 1$ , is zero at the beginning of the execution of the instruction.

[0290] • The first operand length becomes equal to zero during execution of the instruction and normal completion does not occur.

[0291] In one example, the first operand length is zero when the content of general register  $R_1 + 1$  is zero, regardless of the values in the NT and SBB fields of the parameter  
20 block.

[0292] When the first operand length becomes equal to zero during execution of the instruction, the operation ends and the following occurs, in one embodiment:

[0293] • Continuation flag (CF) bit 373 in the parameter block is set to one.

[0294] • Continuation state buffer (CSB) field 392 in the parameter block is  
25 updated.

- [0295]       •       Sub-byte boundary (SBB) field 381 of the parameter block is updated.
- [0296]       •       History length (HL) field 385 of the parameter block is updated.
- [0297]       •       History offset (HO) field 386 of the parameter block is updated, when applicable.
- 5   [0298]       •       Check value field 387 of the parameter block is updated.
- [0299]       •       A model-dependent value is stored to model version number (MVN) field 363 of the parameter block.
- [0300]       •       End-of-block length (EOBL) 389 and end-of block symbol (EOBS) 388 fields of the parameter block are updated.
- 10   [0301]       •       Operation ending supplemental code (OESC) field 365 of the parameter block is set to zeros.
- [0302]       •       The address in general register  $R_1$  is incremented by the number of bytes processed of the first operand that included processing bit 0, and the length in general register  $R_1 + 1$  is decremented by the same number. The number of bytes processed of the first operand that included processing bit 0 is the integer quotient resulting from an integer division with the dividend being the sum of the number of output bits processed and the original value of the SBB, and the divisor being a value of eight.
- 15   [0303]       •       The address in general register  $R_2$  is incremented by the number of source bytes processed, and the length in general register  $R_2 + 1$  is decremented by the same number.
- 20   [0304]       •       Condition code 1 is set.
- [0305]       The formation and updating of the addresses and lengths are dependent on the addressing mode.
- [0306]       When the first operand length is zero at the beginning of the execution of the instruction, the operation ends and the following occurs, in one embodiment:
- 25   [0306]       When the first operand length is zero at the beginning of the execution of the instruction, the operation ends and the following occurs, in one embodiment:

[0307]       •       Condition code 1 is set.

[0308]       Subsequent to the instruction ending with condition code 1 set, it is expected that the program modifies the first operand length, first operand address, or both and re-executes the instruction to resume the operation.

5   [0309]       A PER storage alteration event is recognized, when applicable, for the following:

[0310]       •       Stores to the parameter block, as described below.

[0311]       •       Stores to the first operand location.

[0312]       •       Stores to the third operand location, which occur, e.g., when the history  
10   buffer type (HBT) is one (circular).

[0313]       When the entire parameter block overlaps the PER storage area designation, a PER storage alteration event is recognized, when applicable, for the parameter block. When only a portion of the parameter block overlaps the PER storage area designation, it is model-dependent which of the following occurs:

15   [0314]       •       A PER storage alteration event is recognized when applicable, for the parameter block.

[0315]       •       A PER storage alteration event is recognized, when applicable, for the portion of the parameter block that is stored.

[0316]       A PER zero-address-detection event is recognized when applicable, for the  
20   parameter block, first operand location, second operand location, and third operand location when the HBT is one (circular).

[0317]       Condition code 2 is not applicable to the DFLTC-CCMPR function.

[0318]       When the instruction ends with condition code 1 or 3 set, input data referenced from the second operand location may be completely, or only partially, processed. When  
25   input data is only partially processed, results in the first operand location, first operand

address, first operand length, and SBB field of the parameter block do not represent a state consistent with the updated second operand address and length. In these cases, partially processed data and internal-state information may be placed in the CSB field of the parameter block. The amount of partially processed data depends on conditions existing at the time the operation ends and the model. Although some data may only be partially processed, results stored to the left of the location designated by the updated first operand address are complete and will not be modified when the operation resumes. Furthermore, it is expected that the program subsequently re-executes the instruction to resume the operation, at which time the contents of the CSB field are referenced prior to resuming the operation. When the instruction ends with condition code 0 set, all data is completely processed and all results associated with input and output data represent a consistent state.

[0319] Subsequent to the instruction ending with a non-zero condition code set, and prior to re-executing the instruction for the purpose of resuming the operation, the program should not modify any fields of the parameter block; otherwise results are unpredictable.

[0320] Function Code 4: DFLTCC-XPND (Expand)

[0321] When the DFLTCC-XPND function is specified, an uncompressing operation is performed. The operation includes decoding compressed data symbols from the second operand location into uncompressed data, which is stored to the first operand location.

[0322] In one example, the DFLTCC-XPND function uses a parameter block, an example of which is described above with respect to FIGS. 3K-3L.

[0323] One example of the DFLTCC-XPND operation is described below with respect to Uncompressing Data.

[0324] Normal completion occurs when all elements of the final block of the data set in the second operand are decoded and all uncompressed data is stored to the first operand location. The last block of the data set is identified when the BFINAL bit of the block header is one. When the operation ends due to normal completion, the following occurs, in one embodiment:

- [0325]       •       A model-dependent value is stored to model version number (MVN) field 363 of the parameter block.
- [0326]       •       Continuation flag (CF) field 373 of the parameter block is set to zero.
- [0327]       •       Sub-byte boundary (SBB) field 381 of the parameter block is updated.
- 5   [0328]       •       History length (HL) field 385 of the parameter block is updated.
- [0329]       •       History offset (HO) field 386 of the parameter block is updated, when applicable.
- [0330]       •       Compressed dynamic-Huffman table (CDHT) 367 and compressed dynamic-Huffman table length (CDHTL) fields 366 of the parameter block are set to zeros.
- 10   [0331]       •       Operation ending supplemental code (OESC) field 365 of the parameter block is set to zeros.
- [0332]       •       Check value field 387 of the parameter block is updated.
- [0333]       •       The address in general register R1 is incremented by the number of bytes stored at the first operand location, and the length in general register R1 + 1 is
- 15   decremented by the same number.
- [0334]       •       The address in general register R2 is incremented by the number of bytes processed of the second operand that included processing bit 0, and the length in general register R2 + 1 is decremented by the same number. The number of bytes processed of the second operand that included processing bit 0 is the integer quotient resulting from an integer
- 20   division with the dividend being the sum of the number of input bits processed and the original value of the SBB, and the divisor being a value of eight.
- [0335]       •       Condition code 0 is set.
- [0336]       The formation and updating of the addresses and lengths are dependent on the addressing mode.

[0337] When normal completion occurs, CSB field 392 of the parameter block is undefined after the operation ends.

[0338] When a CPU-determined number of bytes have been processed, the operation ends and the following occurs, in one embodiment:

5 [0339] • Continuation flag (CF) bit 373 in the parameter block is set to one.

[0340] • Continuation state buffer (CSB) field 392 in the parameter block is updated.

[0341] • Sub-byte boundary (SBB) field 381 of the parameter block is updated.

[0342] • Compressed dynamic-Huffman table (CDHT) 367 and compressed  
10 dynamic-Huffman table length (CDHTL) 366 fields of the parameter block are updated.

When partial completion occurs while processing a block with BTYPE value of 10 binary, the bytes of the CDHT field not required to represent the table are stored as zeros. When partial completion occurs while processing a block with BTYPE value of 00 or 01 binary, zeros are stored to the CDHT and CDHTL fields.

15 [0343] • History length (HL) field 385 of the parameter block is updated.

[0344] • History offset (HO) field 386 of the parameter block is updated, when applicable.

[0345] • Check value field 387 of the parameter block is updated.

[0346] • A model-dependent value is stored to model version number (MVN)  
20 field 363 of the parameter block.

[0347] • Operation ending supplemental code (OESC) field 365 of the parameter block is set to zeros.

[0348] • Incomplete function status (IFS) field 383 of the parameter block is updated.

[0349] • Incomplete function length (IFL) field 384 of the parameter block is updated, when applicable.

[0350] • The address in general register  $R_1$  is incremented by the number of bytes stored at the first operand location, and the length in general register  $R_1 + 1$  is decremented by the same number.

[0351] • The address in general register  $R_2$  is incremented by the number of bytes processed of the second operand that included processing bit 0, and the length in general register  $R_2 + 1$  is decremented by the same number. The number of bytes processed of the second operand that included processing bit 0 is the integer quotient resulting from an integer division with the dividend being the sum of the number of input bits processed and the original value of the SBB, and the divisor being a value of eight.

[0352] • Condition code 3 is set.

[0353] The formation and updating of the addresses and lengths are dependent on the addressing mode.

15 [0354] The CPU-determined number of bytes depends on the model, and may be a different number each time the instruction is executed.

[0355] Subsequent to the instruction ending with condition code 3 set, it is expected the program does not modify any input or output specification for the instruction and branches back to re-execute the instruction to resume the operation.

20 [0356] In certain situations, despite ending the instruction with condition code 3 set, the parameter block and general registers are not updated. These situations may occur when the CPU performs a quiescing operation or CPU retry while executing the DEFLATE Conversion Call instruction. In these cases, the CPU-determined number of bytes processed is zero, data may have been stored to the first operand location, data may have been stored to the third operand location, when applicable, and corresponding change bits have been set.

25 [0357] The second operand length is insufficient to complete the operation when the following applies, for instance:

[0358]       •       The last element of a compressed data block with BFINAL equal to one has not been decoded during the operation, and the number of bits in the second operand, as designated by the second operand length and SBB, is less than the number of bits of the next element to decode and all results from decoding data from the second operand location  
5   have been placed at the first operand location.

[0359]       When the second operand length is insufficient to complete the operation, the operation has been partially completed, the operation ends, and the following occurs, in one embodiment:

[0360]       •       Continuation flag (CF) bit 373 in the parameter block is set to one.

10   [0361]       •       Continuation state buffer (CSB) field 392 in the parameter block is updated.

[0362]       •       Sub-byte boundary (SBB) field 381 of the parameter block is updated.

[0363]       •       Compressed dynamic-Huffman table (CDHT) 367 and compressed dynamic-Huffman-table length (CDHTL) fields 366 of the parameter block are updated.

15   When partial completion occurs while processing a block with BTYPE value of 10 binary, the bytes of the CDHT field not required to represent the table are stored as zeros. When partial completion occurs while processing a block with BTYPE value of 00 or 01 binary, zeros are stored to the CDHT and CDHTL fields.

[0364]       •       History length (HL) field 385 of the parameter block is updated.

20   [0365]       •       History offset (HO) field 386 of the parameter block is updated, when applicable.

[0366]       •       Check value field 387 of the parameter block is updated.

[0367]       •       A model-dependent value is stored to the model version number (MVN) field 363 of the parameter block.

- [0368]        •        Operation ending supplemental code (OESC) field 365 of the parameter block is set to zeros.
- [0369]        •        Incomplete function status (IFS) field 383 of the parameter block is updated.
- 5    [0370]        •        Incomplete function length (IFL) field 384 of the parameter block is updated, when applicable.
- [0371]        •        The address in general register R1 is incremented by the number of bytes stored at the first operand location, and the length in general register R1 + 1 is decremented by the same number.
- 10   [0372]        •        The address in general register R2 is incremented by the number of bytes processed of the second operand that included processing bit 0, and the length in general register R2 + 1 is decremented by the same number. The number of bytes processed of the second operand that included processing bit 0 is the integer quotient resulting from an integer division with the dividend being the sum of the number of input bits processed and the
- 15   original value of the SBB, and the divisor being a value of eight.
- [0373]        •        Condition code 2 is set.
- [0374]        The formation and updating of the addresses and lengths are dependent on the addressing mode.
- [0375]        Subsequent to the instruction ending with condition code 2 set, it is expected
- 20   that the program modifies the second operand length, second operand address, or both and re-executes the instruction to resume the operation.
- [0376]        The first operand length is insufficient to complete the operation when the following applies, for instance:
- [0377]        •        Results from decoding data from the second operand location cannot be
- 25   placed at the first operand location due to the first operand length being equal to zero.

[0378] When the first operand length is insufficient to complete the operation, the operation has been partially completed, the operation ends, and the following occurs, in one embodiment:

[0379] • Continuation flag (CF) bit 373 in the parameter block is set to one.

5 [0380] • Continuation state buffer (CSB) field 392 in the parameter block is updated.

[0381] • Sub-byte boundary (SBB) field 381 of the parameter block is updated.

[0382] • The compressed dynamic-Huffman table (CDHT) 367 and compressed dynamic-Huffman-table length (CDHTL) fields 366 of the parameter block are updated.

10 When partial completion occurs while processing a block with BTYPE value of 10 binary, the bytes of the CDHT field not required to represent the table are stored as zeros. When partial completion occurs while processing a block with BTYPE value of 00 or 01 binary, zeros are stored to the CDHT and CDHTL fields.

[0383] • History length (HL) field 385 of the parameter block is updated.

15 [0384] • History offset (HO) field 386 of the parameter block is updated, when applicable.

[0385] • Check value field 387 of the parameter block is updated.

[0386] • A model-dependent value is stored to model version number (MVN) field 363 of the parameter block.

20 [0387] • Operation ending supplemental code (OESC) field 365 of the parameter block is set to zeros.

[0388] • Incomplete function status (IFS) field 383 of the parameter block is updated.

25 [0389] • Incomplete function length (IFL) field 384 of the parameter block is updated, when applicable.

[0390]       •       The address in general register  $R_1$  is incremented by the number of bytes stored at the first operand location, and the length in general register  $R_1 + 1$  is decremented by the same number.

5   [0391]       •       The address in general register  $R_2$  is incremented by the number of bytes processed of the second operand that included processing bit 0, and the length in general register  $R_2 + 1$  is decremented by the same number. The number of bytes processed of the second operand that included processing bit 0 is the integer quotient resulting from an integer division with the dividend being the sum of the number of input bits processed and the original value of the SBB, and the divisor being a value of eight.

10   [0392]       •       Condition code 1 is set.

[0393]       The formation and updating of the addresses and lengths are dependent on the addressing mode.

15   [0394]       Subsequent to the instruction ending with condition code 1 set, it is expected that the program modifies the first operand length, first operand address, or both and re-executes the instruction to resume the operation.

[0395]       A PER storage-alteration event is recognized, when applicable, for the following:

[0396]       •       Stores to the parameter block, as described herein.

[0397]       •       Stores to the first operand location.

20   [0398]       •       Stores to the third operand location, which occur, e.g., when the history buffer type (HBT) is one (circular).

[0399]       In one example, when the entire parameter block overlaps the PER storage area designation, a PER storage alteration event is recognized, when applicable, for the parameter block. When only a portion of the parameter block overlaps the PER storage area designation, it is model-dependent which of the following occurs, in one embodiment:

25

[0400] • A PER storage alteration event is recognized when applicable, for the parameter block.

[0401] • A PER storage alteration event is recognized, when applicable, for the portion of the parameter block that is stored.

- 5 [0402] A PER zero-address-detection event is recognized when applicable, for the parameter block, first operand location, second operand location, and third operand location when the HBT is one (circular).

[0403] When the instruction ends with condition code 1, 2, or 3 set, input data referenced from the second operand location may be completely, or only partially, processed.

- 10 When input data is only partially processed, results in the first operand location, first operand address, first operand length, SBB field of the parameter block, check value field of the parameter block, HL field of the parameter block, IFS field of the parameter block, and when applicable, the third operand location and HO field of the parameter block, do not represent a state consistent with the updated second operand address and length. In these cases, partially
- 15 processed data and internal state information may be placed in the CSB field of the parameter block. The amount of partially processed data depends on conditions existing at the time the operation ends and the model. Although some data may only be partially processed, results stored to the left of the location designated by the updated first operand address are complete and will not be modified when the operation resumes. Furthermore, it is expected that the
- 20 program subsequently re-executes the instruction to resume the operation, at which time the contents of the CSB field are referenced prior to resuming the operation. When the operation ends with condition code 0 set, all data is completely processed and all results associated with input and output data represent a consistent state.

- [0404] Subsequent to the instruction ending with a non-zero condition code set, and
- 25 prior to re-executing the instruction for the purpose of resuming the operation, the program should not modify any fields of the parameter block; otherwise results are unpredictable.

[0405] Compressed Data Blocks

[0406] In one example, the bytes of a compressed data block in storage are processed from, e.g., left to right. Compressed data blocks may, or may not, start or end on byte boundaries. A compressed data block is, for instance, a bit stream. Elements of the block are loaded into storage one bit at a time. The bit stream is loaded, e.g., from right to left within  
 5 each byte of storage and in byte order from, e.g., left to right. When the element is a Huffman code, the bits are stored in order from, e.g., most significant bit to least significant bit of the element. When the element is not a Huffman code, the bits are stored in order from, e.g., least significant bit to most significant bit of the element.

[0407] FIG. 6 illustrates an example of a block 600 with block type 00 binary, which  
 10 contains no compressed data symbols. The following applies to this example, in one embodiment:

[0408] • Compressed data block 600 consists of a bit stream 602 which begins with bit 4 of byte 0, identified as b0, and ends with bit 0 of byte 7, identified as b60.

[0409] • The first element encountered in the bit stream is BFINAL (block  
 15 header final bit) in bit 4 of byte 0.

[0410] • The second element encountered in the bit stream is BTYPE (block type) in bits 2-3 of byte 0. In this example, the BTYPE is 00 binary.

[0411] • Bits to the left of the BTYPE and to the right of a byte boundary are ignored when the BTYPE is 00 binary, which is bits 0-1 of byte 0 in this example.

20 [0412] • The third element encountered in the bit stream is the least significant byte (LSB) of the LEN field, which is followed by the most significant byte (MSB) of the LEN field. The LEN field specifies the number of bytes in the block with literal data. Literal data is, e.g., uncompressed data. The bytes with literal data follow the NLEN field in the bit stream. NLEN is the one's complement of LEN. In one example, bytes 1-2 contain the LEN  
 25 field in little-endian byte order.

[0413] • The elements encountered in the bit stream following the LEN field are the least significant byte of the NLEN field, followed by the most significant byte of the

NLEN field, respectively. Bytes 3-4 contain the NLEN field in little-endian byte order. The NLEN field is the one's complement of the LEN field.

[0414]           •       Elements encountered in the bit stream following the NLEN field are uncompressed data, identified as literal bytes. Bytes 5-7 contain uncompressed data, which is  
5       unchanged from the source data used to generate this block.

[0415]           •       None of the elements contained in this block are Huffman codes. Every element in this block is stored to the bit stream order in order from least significant bit to most significant bit of the element, as defined by the DEFLATE standard. Since the LEN, NLEN, and literal elements are each an integral number of bytes aligned on byte boundaries, these  
10       elements may be processed as units of bytes, and not necessarily as units of bits.

[0416]           FIG. 7 illustrates an example of a block 700 with block type 01 binary, which contains compressed data symbols generated using a fixed-Huffman table (FHT). The following applies to this example, in one embodiment:

[0417]           •       Compressed data block 700 consists of a bit stream 702 which begins  
15       with bit 4 of byte 0, identified as b0, and ends with bit 3 of byte 11, identified as b89.

[0418]           •       The first element encountered in the bit stream is BFINAL in bit 4 of byte 0.

[0419]           •       The second element encountered in the bit stream is BTYPE in bits 2-3 of byte 0. In this example, the BTYPE is 01 binary.

20   [0420]           •       The fixed-Huffman table (FHT) is not a component of the block.

[0421]           •       The third element encountered in the bit stream is the first compressed data symbol, which begins in bit 1 of byte 0. A compressed data symbol consists of the following sub-elements which are encountered in the bit stream in the order in which they are listed, in one example:

25   [0422]           1.       A Huffman code of variable length. The most significant bits of the code designate the length of the code. The code is encountered in the bit stream starting with

the most significant bit of the code and ending with the least significant bit of the code. When the code represents a literal value or the end-of-block symbol, the code is the only sub-element of the compressed data symbol. When the code represents a length of a pointer to the history buffer, the code is followed by subsequent sub-elements of the compressed data symbol.

[0423] 2. When applicable, as specified by the DEFLATE standard, extra length bits may follow the Huffman code representing a pointer length. Extra length bits are encountered in the bit stream starting with the least significant bit and ending with the most significant bit of the extra length bits.

10 [0424] 3. The next sub-element encountered in the bit stream is a 5-bit distance code of a pointer to the history buffer. The distance code is encountered in the bit stream starting with, e.g., the most significant bit of the code and ending with the least significant bit of the distance code.

15 [0425] 4. When applicable, as specified by the DEFLATE standard, extra distance bits may follow the distance code. Extra distance bits are encountered in the bit stream starting with the least significant bit and ending with the most significant bit of the extra distance bits.

[0426] • Bits 0-1 of byte 0, all bits of bytes 1 through 9 and bits 2-7 of byte 10 contain bits of compressed data symbols, as an example.

20 [0427] • The last element encountered in the bit stream is a compressed data symbol containing a single sub-element, which is the Huffman code representing the end-of-block (EOB) symbol. The EOB symbol for a block with BTYPE 01 binary is 00000000 binary. In this example, bit 1 of byte 10 contains the most significant bit of the EOB symbol and bit 3 of byte 11 contains the least significant bit of the EOB symbol.

25 [0428] • Bit 3 of byte 11 contains the last bit of the bit stream, which is the last bit of the compressed data block.

[0429] FIG. 8 illustrates an example of a block 800 with block type 10 binary, which contains compressed data symbols generated using a dynamic-Huffman table (DHT). The following applies to this example, in one embodiment:

[0430] • Compressed data block 800 consists of a bit stream 802 which begins  
5 with bit 4 of byte 0, identified as b0, and ends with bit 3 of byte 11, identified as b89.

[0431] • The first element encountered in the bit stream is BFINAL in bit 4 of byte 0.

[0432] • The second element encountered in the bit stream is BTYPE in bits 2-3 of byte 0. In this example, the BTYPE is 10 binary.

10 [0433] • The third element encountered in the bit stream is the compressed representation of the dynamic-Huffman table (DHT), which begins in bit 1 of byte 0. The compressed representation of the DHT consists of the following sub-elements, which are encountered in the bit stream in the order which they are listed, in one example:

[0434] 1. HLIT: The sum of the 5-bit HLIT sub-element and 257 specifies the  
15 number of Huffman codes representing literal bytes, an EOB symbol, and duplicate string lengths. Valid values of HLIT range from, e.g., 0 to 29. HLIT bits are encountered in the bit stream starting with the least significant bit and ending with the most significant bit of the HLIT sub-element. In this example, bit 1 of byte 0, identified as b3 is the least significant bit of the HLIT sub-element.

20 [0435] 2. HDIST: The sum of the 5-bit HDIST sub-element and 1 specifies the number of Huffman codes representing duplicate string pointer distances. Valid values of HDIST range from, e.g., 0 to 29. HDIST bits are encountered in the bit stream starting with the least significant bit and ending with the most significant bit of the HDIST sub-element.

[0436] 3. HCLLEN: The sum of the 4-bit HCLLEN sub-element and 4 specifies the  
25 number of Huffman codes representing code lengths. Valid values of HCLLEN range from, e.g., 0 to 15. HCLLEN bits are encountered in the bit stream starting with the least significant bit and ending with the most significant bit of the HCLLEN sub-element.

[0437] 4. A sequence of codes specifying a bit length for each of the code lengths defined for the compressed DHT. The number of codes is equal to the sum of HCLen and 4. Each code is 3 bits.

[0438] 5. A sequence of codes specifying a code length for each of the elements of the set consisting of literal bytes, an EOB symbol, and duplicate string lengths. The number of code lengths specified is equal to the sum of HLIT and 257.

[0439] When the last code length (CL) for the set of literal bytes, an EOB symbol, and duplicate string lengths is 16, 17, or 18, and the extra bits following the CL specify repeating the CL for more elements than are defined for the set, the code length also applies to the set of duplicate string pointer distances. The sequence of codes specifying code lengths for the set of literal bytes, an EOB symbol and duplicate string lengths, followed by the sequence of codes specifying code lengths for duplicate string pointer distances is a contiguous sequence for both sets.

[0440] 6. A sequence of codes specifying a code length for each of the elements of the set consisting of duplicate string pointer distances. The number of code lengths specified is equal to the sum of HDIST and 1.

[0441] • The fourth element encountered in the bit stream is the first compressed data symbol. A compressed data symbol consists of the following sub-elements, which are encountered in the bit stream in the order which they are listed, in one embodiment:

[0442] 1. A Huffman code of variable length. The most significant bits of the code designate the length of the code. The code is encountered in the bit stream starting with the most significant bit of the code and ending with the least significant bit of the code. When the code represents a literal value or the end-of-block symbol, the code is the only sub-element of the compressed data symbol. When the code represents a length of a pointer to the history buffer, the code is followed by subsequent sub-elements of the compressed data symbol.

[0443] 2. When applicable, as specified by the DEFLATE standard, extra length bits may follow the Huffman code representing a pointer length. Extra length bits are

encountered in the bit stream starting with, e.g., the least significant bit and ending with the most significant bit of the extra length bits.

[0444]           3.       The next sub-element encountered in the bit stream is a 5-bit distance code of a pointer to the history buffer. The distance code is encountered in the bit stream  
5   starting with, e.g., the most significant bit of the code and ending with the least significant bit of the distance code.

[0445]           4.       When applicable, as specified by the DEFLATE standard, extra distance bits may follow the distance code. Extra distance bits are encountered in the bit stream starting with, e.g., the least significant bit and ending with the most significant bit of  
10   the extra distance bits.

[0446]           •       Subsequent bits encountered in the bit stream, up to and including, e.g., bit 5 of byte 10, contain bits of compressed data symbols.

[0447]           •       The last element encountered in the bit stream is a compressed data symbol containing a single sub-element, which is the Huffman code representing the end-of-  
15   block (EOB) symbol. In this example, bit 4 of byte 10 contains the most significant bit of the EOB symbol and bit 3 of byte 11 contains the least significant bit of the EOB symbol.

[0448]           •       Bit 3 of byte 11 contains the last bit of the bit stream, which is the last bit of the compressed data block.

[0449]           In the above descriptions of the various block types, certain constant values are  
20   specified, as well as specific bits, bytes, directions, etc. These are only examples. Other constant values, bits, bytes, directions, etc. may be specified in other embodiments.

[0450]           Processing a Compressed Data Set

[0451]           Examples of processing a compressed data set to illustrate example uses of the DEFLATE Conversion Call instruction and augment the descriptions of various fields of the  
25   parameter block are provided. The examples do not describe all possible scenarios, requirements, and capabilities, but illustrate various of the scenarios, requirements and/or capabilities. The examples and descriptions apply, for instance, to a compressed data set in

storage, an example of which is illustrated in FIG. 9. As shown, a compressed data set 900 includes a plurality of compressed data blocks 902, and a beginning of data set 900 is indicated by a compressed data set begin address (CDSBA) 904.

[0452] For the examples described herein, it is intended for a program processing the  
5 compressed data set to consider the following, in one embodiment:

[0453] • A single parameter block may be defined and referenced by multiple  
usages of the DEFLATE Conversion Call instruction to process the entire compressed data  
set. Check value 387 and check value type 375 fields of the parameter block shall apply to  
the compressed data blocks (e.g., all blocks) in the compressed data set. Sub-byte boundary  
10 field 381 of the parameter block shall apply to transitions between individual blocks. History  
length 385 and history offset 386 may apply to multiple blocks. The remaining fields of the  
parameter block only apply, in one example, to the individual compressed data block being  
processed by a specific execution of a DEFLATE Conversion Call instruction.

[0454] • An individual check value applies, e.g., to all of the uncompressed data  
15 represented by the compressed data set.

[0455] • There is no history for the first compressed data symbol in block 1 to  
reference. Subsequent symbols in block 1 may reference history corresponding to previously  
encountered symbols in block 1. Symbols in block 2 may reference history corresponding to  
previously encountered symbols in blocks 2 and 1. Symbols in block 3 may reference history  
20 corresponding to previously encountered symbols in blocks 3, 2, and 1.

[0456] FIG. 10 lists one example of a portion of a sample program 1000 used to  
compress data in compressed data set 900 described in FIG. 9. Further, FIG. 11 lists the  
values for certain fields of the parameter block used during execution of the DFLTCC  
instruction located at the instruction address labeled IABLK1 (1002) in FIG. 10. For  
25 instance, FIG. 11 depicts various parameter block fields 1100; values for those fields at the  
start of the compress operation 1102; values for those fields at the end of the operation when  
condition code 1, 2 or 3 is set 1104; and values for those fields at the end of the operation  
when condition code 0 is set 1106.

[0457] Similarly, FIG. 12 lists the values for certain fields of the parameter block used during execution of the DFLTCC instruction located at the instruction address labeled IABLK2 (1004) of FIG. 10. These figures demonstrate some of the details associated with using the DEFLATE Conversion Call instruction multiple times to process an entire  
5 compressed data set.

[0458] Moreover, referring to FIG. 13, one example of a portion of a sample program 1300 used to decompress data from the compressed data set of FIG. 9 is depicted.

[0459] Compressing Data

[0460] The process of compressing data includes generating one or more compressed  
10 data blocks. The compress function of the DEFLATE Conversion Call instruction is used to construct a portion of an individual block. The portion may be the entire block. This function generates portions of a block with block type (BTYPE) 01 or 10 binary, and not 00 binary. When the new task bit (NT) of the parameter block is one, the first block of compressed data is generated and there is no history to reference from previously performed compressing  
15 operations.

[0461] In one example, an individual block contains the following elements in the order which they are listed:

[0462] 1. Final block indication (BFINAL).

[0463] 2. Block type (BTYPE).

20 [0464] 3. Compressed format of a dynamic-Huffman table, when applicable.

[0465] 4. Compressed data symbols.

[0466] 5. End-of-block (EOB) symbol.

[0467] The compression operation generates the elements specified in the order defined for a block. The elements may begin or end between byte boundaries in storage. The  
25 sub-byte boundary (SBB) applies to storing of the first element to the first operand location.

A compressed data block is a bit stream. Components of the block are loaded into storage one bit at a time. As an example, the bit stream is loaded from right to left within each byte of storage and in byte order from left to right.

[0468] When the SBB is non-zero, the reference to the first byte at the first operand  
5 location is an update reference.

[0469] Uncompressed data from the second operand location is compressed and stored as compressed data symbols to the first operand location.

[0470] When the first operand length is zero at the beginning of the execution of the instruction, the first operand is not accessed, and the first operand address and first operand  
10 length in general registers R1 and R1 + 1, respectively, are not changed. This applies when the value of CF field 373 (FIG. 3L) is zero or one at the beginning of execution of the instruction.

[0471] When the second operand length is zero at the beginning of execution of the instruction, the second operand is not accessed, and the second operand address and second  
15 operand length in general registers R2 and R2 + 1, respectively, are not changed. The second operand length is zero at the beginning of execution of the instruction for the following case, as an example:

[0472] • The instruction is being re-executed to resume the operation (CF field 373 of the parameter block is one at the beginning of execution of the instruction) and  
20 completing the operation can be performed with references to CSB field 392 of the parameter block, and without references to the second operand.

[0473] In one embodiment, the program is not to use the DEFLATE Conversion Call instruction to perform the following operations:

[0474] • Generate an empty compressed data block. An empty compressed data  
25 block consists of, e.g., a block header, a compressed format of a DHT when applicable, and an EOB symbol.

[0475]       •       Close an open compressed data block. That is, only store an EOB symbol to the end of the compressed data block.

[0476]       The compression algorithm includes searching an updated history of recently compressed data for a string of bytes which matches data currently being compressed from  
5   the second operand location. Before the compression operation begins or resumes, the following applies, in one embodiment:

[0477]       •       When new task (NT) 374 is one, there is no initial history available to reference.

[0478]       •       When NT is zero, and bit 56 of general register 0 (HBT) is zero (in-  
10   line), the initial history available to reference is located to the left of, and adjacent to, the leftmost byte of the second operand, and the length of the initial history is specified by history length (HL) field 385 of the parameter block.

[0479]       •       When NT is zero, and bit 56 of general register 0 (HBT) is one  
15   (circular), the initial history available to reference is located in the third operand location, as specified by history offset (HO) 386 and history length (HL) 385 fields of the parameter block.

[0480]       During the compression operation, fetch-type references to the entire history may be made, regardless of which bytes of history are used to perform the operation.  
Furthermore, when the history buffer type is circular, fetch-type references to the entire 32 K-  
20   byte history buffer may be made, regardless of which bytes of history are used to perform the operation.

[0481]       During the compression operation, history is updated. Subsequent to encoding one or more bytes of source data into a compressed data symbol without encountering a general operand data exception condition, the source bytes are concatenated to the end of the  
25   history. The most recently processed bytes of source data, up to a maximum of 32 K-bytes, constitute the updated history available to reference while processing subsequent bytes of source data.

[0482] When the compression operation ends, the following applies, in one example, to the resulting history available to subsequently resume the operation, or begin another operation:

[0483] • When the HBT is in-line, storage updates to the second operand

5 location are not required when the history is updated. The updated second operand address and updated HL specify the updated location and updated length of the resulting history.

[0484] • When the HBT is circular, storage updates to the third operand location are performed when the history is updated. The third operand address, updated HO, and updated HL specify the updated location and updated length of the resulting history.

10 [0485] As examples, FIGS. 14A-14C illustrate the location of an in-line history buffer with respect to the second operand before and after multiple executions of a DEFLATE Conversion Call instruction with the DFLTCC-CMPR function specified, as well as in-line history specified (e.g., bit 310 = 0), when each execution ends with partial completion. For instance, FIG. 14A depicts in-line history before DFLTCC-CMPR execution number 1; FIG.  
15 14B depicts in-line history before DFLTCC-CMPR execution number 2 and after execution number 1; and FIG. 14C depicts in-line history after DFLTCC-CMPR execution number 2. The explanation provided in FIG. 14C also applies to FIGS. 14A and 14B.

[0486] When the HBT (history buffer type) specified by bit 56 of general register 0 is circular (e.g., bit 310 = 1), the history is maintained in, for instance, a 32 K-byte buffer  
20 located at the third operand location. The location of the first byte of history within the buffer (HB) is designated by, e.g., the sum of the contents of general register R3 and history offset (HO) 386 (FIG. 3L). The first byte of history is the least recently processed byte of uncompressed data in the buffer. The location of the last byte of history within the buffer (HE) is designated by the following equation, as an example:

25 [0487] 
$$HE = R_3 + \text{modulo}32K(HO + HL - 1)$$

[0488] The last byte of history is the most recently processed byte of uncompressed data in the buffer. When the sum of history offset (HO) 386 (FIG. 3L) and history length (HL)

385 exceeds the size of the third operand (e.g., 32 K-bytes), the history wraps from the end of the third operand to the beginning of the third operand.

[0489] As examples, FIGS. 15A-15E illustrate the location of the history within a circular history buffer before and after multiple executions of a DEFLATE Conversion Call instruction with the DFLTCC-CMPR function specified, as well as a circular history buffer specified (bit 310 = 1), when each execution ends with partial completion. For instance, FIG. 15A depicts a circular history buffer before DFLTCC execution number 1; FIG. 15B depicts a circular buffer before DFLTCC execution number 2 and after execution number 1; FIG. 15C depicts a circular buffer before DFLTCC execution number 3 and after execution number 2; FIG. 15D depicts the circular buffer before DFLTCC execution number 4 and after execution number 3; and FIG. 15E depicts the circular buffer after DFLTCC execution number 4. The explanation provided in FIG. 15E also applies to FIGS. 15A-15D.

[0490] When the HBT is circular and the number of bytes processed from the second operand location is less than, e.g., 32,768, the following applies, in one example:

15 [0491] • Stores are made to a range of bytes in the third operand location. The range of bytes includes and starts with the location designated by, for instance:

[0492]  $R_3 + \text{modulo}32K(HOO + HLO)$ , where

[0493] HOO: The history offset before the instruction executes.

[0494] HLO: The history length before the instruction executes.

20 [0495] The range of bytes includes and ends with the location designated by, e.g.:

[0496]  $R_3 + \text{modulo}32K(HOO + HLO + BP - 1)$ , where

[0497] BP: The number of bytes processed from the second operand location during execution of the instruction.

[0498] Stores made to the range of bytes just described are subject to store-type access exceptions, PER storage alteration events, and setting change bits, as an example.

[0499] • Stores which do not modify the contents of storage locations and are not necessary, may be made to bytes in the third operand location which are not included in the range just described. Stores to such locations are also subject to store type access exceptions, PER storage alteration events, and setting change bits.

5 [0500] When the HBT is circular and the number of bytes processed from the second operand location is more than, or equal to, e.g., 32,768, stores are made to all bytes of the third operand location and subject to store type access exceptions, PER storage alteration events, and setting change bits.

[0501] When the block continuation flag (BCF) 377 is zero, a 3 bit block header, including BFINAL followed by BTYPE, is stored to the first operand location. The BFINAL bit of the block header is set equal to the block header final bit (BHF) 379 of the parameter block. When the Huffman table type (HTT) 376 is zero, the BTYPE field of the block header is set to, e.g., 01 binary and when the HTT is one, the BTYPE field of the block header is set to, e.g., 10 binary. When a block header is stored, the BFINAL bit is stored to the bit  
15 specified by the SBB in the first byte of the first operand. Subsequently, the BTYPE is stored to the first operand location. When the BCF is one, a block header is not stored.

[0502] When the Huffman table type (HTT) is one, the compressed format of the dynamic- Huffman table (DHT) 367 specified in the parameter block is examined for general operand data exception conditions. When a general operand data exception condition exists  
20 for the specified compressed format of the DHT, the compressed DHT is referred to as invalid and is not to be used to compress data. Example definitions of general operand data exception conditions are described further below. When the compressed format of the DHT specifies a bit length for a code length, or a code length for a literal byte, the EOB symbol, a duplicate string length, or a duplicate string pointer distance, which is greater than the length required  
25 by the Huffman algorithm to specify a proper and functional Huffman tree, the compressed DHT is still used to derive a functional DHT and compress data. When the block continuation flag (BCF) is zero and the HTT is one, the compressed format of the DHT, as specified in CDHT field 367 of the parameter block is stored to the first operand location.

[0503] During the compression operation, source data from the second operand location is encoded into compressed data symbols. As part of the encoding, source data is compared to the history. When no match is found, the intermediate representation of the source data is literal bytes, which is the same as the source data. When a match is found, the intermediate representation of the source data is a pointer to a location within the history which contains a duplicate copy of the source data. A pointer consists of a length and a distance. The length is the number of source data bytes which match a string in the history. The distance is the number of bytes from the end of the history to the beginning of the string which matches the source data. Two Huffman code trees from the Huffman table are used, in one example, to encode the intermediate representation of the source data into compressed data symbols. When the Huffman table type (HTT) is zero, a fixed-Huffman table (FHT), as described by the DEFLATE standard, specifies the two Huffman code trees used for encoding intermediate results. When HTT 376 is one, the dynamic-Huffman table (DHT), which is derived from the compressed representation of the DHT, specified in CDHT field 367 of the parameter block, specifies the two Huffman code trees used for encoding intermediate results. The encoding is performed as described by the DEFLATE standard. When a non-universal DHT is used which does not specify a Huffman code to be used to encode the intermediate representation of the source data, a general operand data exception is recognized. The bits of the resulting compressed data symbol are arranged in the order specified by the DEFLATE standard before storing the result to the first operand location.

[0504] In one example, duplicate-string lengths range from 3 to 258 bytes.

[0505] Prior to processing further source data, the history is updated, as described herein.

[0506] The process is repeated, in one example, until all source bytes have been processed.

[0507] After the source bytes (e.g., all source bytes) have been processed and the block closing control (BCC) 378 is one, an end-of-block (EOB) symbol is stored to the first operand location. When a fixed-Huffman table is used, Huffman code 0000000 binary is used for the EOB symbol. When a dynamic-Huffman table (DHT) is used, the Huffman code used

for the EOB symbol is specified by the DHT. The bits of the EOB symbol are arranged in the order specified by the DEFLATE standard before storing the EOB symbol to the first operand location.

[0508] When the last compressed data symbol of the operation (including the EOB symbol), only occupies a portion of the last byte to store, the bits that do not contain a portion of the last symbol are stored as zeros, in one example.

[0509] Subsequent to processing the last compressed data symbol, the following occurs, in one embodiment:

[0510] • A model-dependent value is stored to model version number (MVN) field 363 of the parameter block.

[0511] • Sub-byte boundary (SBB) field 381 of the parameter block is updated.

[0512] • End-of-block length (EOBL) 389 and end-of block symbol (EOBS) 388 fields of the parameter block are updated.

[0513] • The address in general register R1 is incremented by the number of bytes processed of the first operand that included processing bit 0, and the length in general register R1 + 1 is decremented by the same number. The number of bytes processed of the first operand that included processing bit 0 is the integer quotient resulting from an integer division with the dividend being the sum of the number of output bits processed and the original value of the SBB, and the divisor being a value of eight.

[0514] • The address in general register R<sub>2</sub> is incremented by the number of source bytes processed, and the length in general register R<sub>2</sub> + 1 is decremented by the same number.

[0515] The formation and updating of the addresses and lengths are dependent on the addressing mode.

[0516] Coincident with compressing the source data, the source data is an input to generating a 32-bit check value, described above. The resulting check value is stored to check value field 387 of the parameter block.

[0517] Uncompressing Data

5 [0518] In one embodiment, the expand function of the DEFLATE Conversion Call instruction is used to decode a compressed data set into uncompressed data. The compressed data set in the second operand location includes one or more consecutive compressed data blocks. The blocks of the data set are processed from left to right, in one example, and the bytes of a block are processed, e.g., from left to right. The blocks may or may not start or end  
10 on byte boundaries. Each block is decoded independent of other blocks in the data set. General register R<sub>2</sub> specifies the logical address of the leftmost byte of the first block in the data set. The last block in the data set is the block encountered during processing with the BFINAL bit equal to one. In one example, there are three types of blocks to process. The technique of decoding the contents of a block is a function of the block type (BTYPED).

15 [0519] When the operation begins (e.g., when continuation flag field 373 of the parameter block is zero), the bit designated by general register R<sub>2</sub>, new task (NT) field 374, and sub-byte boundary (SBB) field 381 is interpreted as the first bit of a compressed data block (the BFINAL bit of a block header).

[0520] The expand function includes referencing an updated history of recently  
20 decoded uncompressed data. Before the uncompressing operation begins or resumes, the following applies, in one embodiment:

[0521] • When new task (NT) 374 is one, there is no initial history available to reference.

[0522] • When NT is zero, and bit 56 of general register 0 (HBT) is zero (in-  
25 line), the initial history available to reference is located to the left of, and adjacent to, the leftmost byte of the first operand, and the length of the initial history is specified by history length (HL) field 385 of the parameter block.

[0523]       •       When NT is zero, and bit 56 of general register 0 (HBT) is one (circular), the initial history available to reference is located in the third operand location, as specified by history offset (HO) 386 and history length (HL) 385 fields of the parameter block.

5   [0524]       During the operation, fetch-type references to the entire history may be made, regardless of which bytes of history are used to perform the operation. Furthermore, when the history buffer type is circular, fetch-type references to the entire history buffer (e.g., 32 K-byte) may be made, regardless of which bytes of history are used to perform the operation.

10 [0525]       During the uncompressing operation, history is updated. Subsequent to decoding source data without encountering a general operand data exception condition, the resulting bytes of uncompressed data are concatenated to the end of the history. The most recently decoded bytes of uncompressed data, up to a maximum of, e.g., 32 K-bytes, constitute the updated history available to reference while processing subsequent source data.

15 [0526]       When the uncompressing operation ends, the following applies to the resulting history available to subsequently resume the operation, or begin another operation, in one example:

[0527]       •       When the HBT is in-line, storage updates to the first operand location also constitute updates to the resulting history. The updated first operand address and updated HL specify the updated location and updated length of the resulting history.

20 [0528]       •       When the HBT is circular, storage updates to the third operand location are performed when the history is updated. The third operand address, updated HO, and updated HL specify the updated location and updated length of the resulting history.

25 [0529]       As examples, FIGS. 16A-16C illustrate examples of the location of an in-line history buffer with respect to the first operand before and after multiple executions of a DEFLATE Conversion Call instruction with the DFLTCC-XPND function specified, as well as in-line history specified, when each execution ends with partial completion. History length (HL) 385 is modified during the operation. For instance, FIG. 16A depicts one example of the in-line history prior to DFLTCC-XPND execution number 1; FIG. 16B depicts an

example of the in-line history before DFLTCC-XPND execution number 2 and after execution number 1; and FIG. 16C depicts an example of the in-line history after DFLTCC-XPND execution number 2. The explanation provided in FIG. 16C also applies to FIGS. 16A-16B.

- 5 [0530] When the HBT specified by bit 56 of general register 0 is circular, the history is maintained in, for instance, a 32 K-byte buffer located at the third operand location. The location of the first byte of history within the buffer (HB) is designated by the sum of the contents of general register R3 and history offset (HO) 386. The first byte of history is the least recently processed byte of uncompressed data in the buffer. The location of the last byte  
10 of history within the buffer (HE) is designated by, e.g., the following equation:

$$[0531] \quad HE = R_3 + \text{modulo}32K(HO + HL - 1).$$

- [0532] The last byte of history is the most recently processed byte of uncompressed data in the buffer. When the sum of the history offset (HO) and history length (HL) exceeds the size of the third operand (e.g., 32 K-bytes), the history wraps from the end of the third  
15 operand to the beginning of the third operand. FIGS. 15A-15E, described herein, illustrate examples of the location of the history within a circular history buffer before and after multiple executions of a DEFLATE Conversion Call instruction with the DFLTCC-XPND function and a circular history buffer specified, when each execution ends with partial completion.

- 20 [0533] When the HBT is circular and the number of bytes stored to the first operand location is less than, e.g., 32,768, the following applies, in one example:

[0534] • Stores are made to a range of bytes in the third operand location. The range of bytes includes and starts with the location designated by:

$$[0535] \quad R_3 + \text{modulo}32K(HOO + HLO), \text{ where}$$

- 25 [0536] HOO: The history offset before the instruction executes.

[0537] HLO: The history length before the instruction executes.

[0538] The range of bytes includes and ends with the location designated by, e.g.,

[0539]  $R_3 + \text{modulo}32K(HOO + HLO + BP - 1)$ , where

[0540] BP: The number of bytes stored to the first operand location during execution of the instruction.

5 [0541] Stores made to the range of bytes just described are subject to store-type access exceptions, PER storage-alteration events, and setting change bits.

[0542] • Stores which do not modify the contents of storage locations and are not necessary, may be made to bytes in the third operand location which are not included in the range just described. Stores to such locations are also subject to store-type access

10 exceptions, PER storage-alteration events, and setting change bits.

[0543] When the HBT is circular and the number of bytes stored to the first operand location is more than, or equal to, e.g., 32,768, stores are made to, e.g., all bytes of the third operand location and subject to store-type access exceptions, PER storage-alteration events, and setting change bits.

15 [0544] When the BTYPE is 00 binary, the block does not contain compressed data. FIG. 6, described herein, illustrates one example of a block with BTYPE equal 00 binary. The LEN field specifies the number of literal bytes in the block. The byte order of the LEN field is little-endian. The LEN field may specify zero literal bytes. The literal bytes of the block are placed at the first operand location. The history is also updated, as previously described, with  
20 each literal byte of the block.

[0545] When the BTYPE is 01 binary, the block contains compressed data symbols that were generated using a fixed-Huffman table (FHT). The FHT is defined by the DEFLATE standard and is not part of the block. FIG. 7, as described herein, illustrates one example of a block with BTYPE equal 01 binary. Subsequent to interpreting the block header,  
25 compressed data symbols are decoded in the order in which they appear in the block. Bytes of the block are processed from, e.g., left to right and bits within each byte of the block are processed from, e.g., right to left. In one example, each symbol is completely processed prior

to processing the next symbol in the block. Each symbol which is not the end-of-block (EOB) symbol represents a literal value or a pointer to a substring previously decoded in the history buffer. A previously decoded substring is also referred to as a duplicate string. In one example, duplicate string lengths range from 3 to 258 bytes. A pointer consists of codes  
5 representing the substring length and the distance from the end of the history to the beginning of the substring. When a symbol represents a substring in the history, the substring is referenced from the history buffer. The uncompressed data resulting from decoding a symbol is placed at the first operand location.

[0546] Prior to processing further source data, the history is updated as previously  
10 described.

[0547] The updated history applies to decoding the next symbol of the block. When the EOB symbol is encountered, processing of the block is complete.

[0548] When the BTYPE is 10 binary, the block contains compressed data symbols that were generated using a dynamic-Huffman table (DHT). A compressed format of the DHT  
15 used is an element of the compressed data block. FIG. 8, described herein, illustrates one example of a block with BTYPE equal 10 binary. Subsequent to interpreting the block header, the compressed format of the DHT provided within the compressed data block is examined for general operand data exception conditions. When a general operand data exception  
20 condition exists for the provided compressed format of the DHT, the compressed format of the DHT is referred to as invalid and is not to be used to uncompress data. When the compressed format of the DHT specifies a bit length for a code length, or a code length for a literal byte, the EOB symbol, a duplicate string length, or a duplicate string pointer distance, which is greater than the length required by the Huffman algorithm to specify a proper and functional Huffman tree, the compressed DHT is still used to derive a functional DHT and  
25 compress data. Subsequent to examining the compressed format of the DHT, compressed data symbols are decoded in the order in which they appear in the block. Bytes of the block are processed, e.g., from left to right and bits within each byte of the block are processed from, e.g., right to left. Each symbol is completely processed, in one example, prior to processing the next symbol in the block. The processing of symbols in a block with BTYPE 10 binary is  
30 the same as previously described for processing symbols in a block with BTYPE 01, except

the former uses the DHT provided to decode symbols, and the latter uses the FHT to decode symbols. When a non-universal DHT is provided which does not specify a Huffman code to be used to decode a compressed data symbol, a general operand data exception is recognized.

[0549] Coincident with uncompressing the second operand, the uncompressed data is  
5 an input to generating a check value (e.g., a 32-bit check value). The resulting check value is stored to check value field 387 of the parameter block.

[0550] Subsequent to processing the last block of the data set, the following occurs, in one embodiment:

[0551] • A model-dependent value is stored to model version number (MVN)  
10 field 363 of the parameter block.

[0552] • Sub-byte boundary (SBB) field 381 of the parameter block is updated.

[0553] • The address in general register  $R_1$  is incremented by the number of bytes stored at the first operand location, and the length in general register  $R_1 + 1$  is decremented by the same number.

15 [0554] • The address in general register  $R_2$  is incremented by the number of bytes processed of the second operand that included processing bit 0, and the length in general register  $R_2 + 1$  is decremented by the same number. The number of bytes processed of the second operand that included processing bit 0 is the integer quotient resulting from an integer division with the dividend being the sum of the number of input bits processed and the  
20 original value of the SBB, and the divisor being a value of eight.

[0555] The formation and updating of the addresses and lengths are dependent on the addressing mode.

[0556] When the first operand length is zero at the beginning of the execution of the instruction, the first operand is not accessed, and the first operand address and first operand  
25 length in general registers  $R_1$  and  $R_1 + 1$ , respectively, are not changed. This applies when the value of CF field 373 is zero or one at the beginning of the execution of the instruction.

[0557] When the second operand length is zero at the beginning of the execution of the instruction, the second operand is not accessed, and the second operand address and second operand length in general registers  $R_2$  and  $R_2 + 1$ , respectively, are not changed. The second operand length is zero at the beginning of the execution of the instruction for the following case, in one embodiment:

[0558] • The instruction is being re-executed (e.g., CF field 373 of the parameter block is one at the beginning of the execution of the instruction) and the entire second operand was processed when the instruction was previously executed.

[0559] The uncompressing operation may end without storing any results to the first operand location, even though data was processed from the second operand location. This occurs when the data processed from the second operand location only contains any of the following compressed data block elements, in an example:

[0560] • A block header.

[0561] • The LEN field of a block with block type 00 binary.

15 [0562] • The NLEN field of a block with block type 00 binary.

[0563] • A compressed format of a dynamic-Huffman table.

[0564] • An end-of-block (EOB) symbol.

[0565] The following conditions apply to executing the DEFLATE Conversion Call instruction, in one or more embodiments:

20 [0566] A general operand data exception is recognized when the DFLTCC-GDHT function is specified and the following condition occurs, in one example:

[0567] • The format of the parameter block, as specified by parameter block version number 362, is not supported by the model.

[0568] A general operand data exception is recognized when the DFLTCC-CMPR function is specified and any of the following conditions occur, in one example:

- [0569]       •       The format of the parameter block, as specified by parameter block version number 362, is not supported by the model.
- [0570]       •       NT 374 is zero and HL 385 is greater than, e.g., 32,768.
- [0571]       •       HTT 376 is one and CDHTL 366 is less than, e.g., 42 or greater than,  
5   e.g., 2283.
- [0572]       •       HTT 376 is one and CDHTL 366 does not equal the length of the compressed format of the DHT specified in CDHT field 367.
- [0573]       •       HTT 376 is one and the HLIT sub-element of the compressed format of the DHT is greater than, e.g., 29 (invalid DHT).
- 10   [0574]       •       HTT 376 is one and the HDIST sub-element of the compressed format of the DHT is greater than, e.g., 29 (invalid DHT).
- [0575]       •       HTT 376 is one and the compressed format of the DHT (contents of CDHT field 367) specifies a code which is in the sequence of codes specifying the bit lengths for the, e.g., 19 possible code lengths defined for a compressed DHT, and is less than the  
15   length required by the Huffman algorithm to specify a functional Huffman tree (invalid DHT).
- [0576]       •       HTT 376 is one and the compressed format of the DHT (contents of CDHT field 367) specifies code length, e.g., 16 (copy previous code length) as the first code length for the set of elements consisting of literal bytes, an EOB symbol, and duplicate string  
20   lengths (invalid DHT).
- [0577]       •       HTT 376 is one and the compressed format of the DHT (contents of CDHT field 367) specifies a code which is in the sequence of codes specifying code lengths for literal bytes, and the code does not match any of the codes determined to represent the set of referenced code lengths, as specified earlier in the compressed DHT (invalid DHT).
- 25   [0578]       •       HTT 376 is one and the compressed format of the DHT (contents of CDHT field 367) specifies a code which assigns code length 0 (CL0) to the EOB symbol. In

this case, the corresponding DHT does not specify a Huffman code to represent an EOB symbol (invalid DHT).

[0579]           •       HTT 376 is one and the compressed format of the DHT (contents of the CDHT field 367) specifies a code which is in the sequence of codes specifying code lengths  
5   for duplicate string lengths and pointer distances, and the code does not match any of the codes determined to represent the set of referenced code lengths, as specified earlier in the compressed DHT (invalid DHT).

[0580]           •       HTT 376 is one and the compressed format of the DHT (contents of the CDHT field 367) specifies a number of code lengths which is greater than the number of  
10   Huffman codes in the DHT, as specified by the sum of the values in the HLIT field, the HDIST field, and, e.g., 258. This is possible with improper uses of code lengths 16, 17, and 18, as examples (invalid DHT).

[0581]           •       HTT 376 is one and the compressed format of the DHT (contents of the CDHT field 367) specifies a code length for the set of literal bytes, EOB symbol, and  
15   duplicate string lengths, which is less than the length required by the Huffman algorithm to specify a functional Huffman tree (invalid DHT).

[0582]           •       HTT 376 is one and the compressed format of the DHT (contents of the CDHT field 367) specifies a code length for the set of duplicate string pointer distances,  
20   which is less than the length required by the Huffman algorithm to specify a functional Huffman tree (invalid DHT).

[0583]           •       The CPU attempts to generate a compressed data symbol to represent a literal byte in the second operand, and the DHT derived from the contents of the CDHT field is non-universal and does not specify a Huffman code corresponding to that literal byte.

[0584]           •       The CPU attempts to generate a compressed data symbol to represent a  
25   duplicate string in the second operand, and the DHT derived from the contents of the CDHT field is non-universal and does not specify a Huffman code corresponding to that duplicate string length or pointer distance.

[0585] A general operand data exception is recognized when, for instance, the DFLTCC-XPND function is specified and any of the following conditions occur, as examples:

- [0586] • The format of the parameter block, as specified by parameter block version number 362, is not supported by the model.
- [0587] • NT 374 is zero and HL 385 is greater than, e.g., 32,768.
- [0588] • A compressed data block with BTYPE equal 11 binary is encountered.
- [0589] • A compressed data block with BTYPE equal 00 binary and NLEN not equal to the one's complement of LEN is encountered.
- 10 [0590] • A compressed format of a DHT (contents of a compressed data block with BTYPE equal 10 binary) is encountered and the HLIT sub-element of the compressed DHT is greater than, e.g., 29 (invalid DHT).
- [0591] • A compressed format of a DHT (contents of a compressed data block with BTYPE equal 10 binary) is encountered and the HDIST sub-element of the compressed  
15 DHT is greater than, e.g., 29 (invalid DHT).
- [0592] • A compressed format of a DHT (contents of a compressed data block with BTYPE equal 10 binary) is encountered which specifies a code which is in the sequence of codes specifying the bit lengths for, e.g., the 19 possible code lengths defined for a compressed DHT, and is less than the length required by the Huffman algorithm to specify a  
20 functional Huffman tree (invalid DHT).
- [0593] • A compressed format of a DHT (contents of a compressed data block with BTYPE equal 10 binary) is encountered which specifies code length, e.g., 16 (copy previous code length) as the first code length for the set of elements consisting of literal bytes, an EOB symbol, and duplicate string lengths (invalid DHT).
- 25 [0594] • A compressed format of a DHT (contents of a compressed data block with BTYPE equal 10 binary) is encountered which specifies a code which is in the sequence

of codes specifying code lengths for literal bytes, and the code does not match any of the codes determined to represent the set of referenced code lengths, as specified earlier in the compressed DHT (invalid DHT).

[0595]           •       A compressed format of a DHT (contents of a compressed data block  
5   with BTYPE equal 10 binary) is encountered which specifies a code which assigns code  
length 0 (CL0) to the EOB symbol. In this case, the corresponding DHT does not specify a  
Huffman code to represent an EOB symbol (invalid DHT).

[0596]           •       A compressed format of a DHT (contents of a compressed data block  
10   with BTYPE equal 10 binary) is encountered which specifies a code which is in the sequence  
of codes specifying code lengths for duplicate string lengths and pointer distances, and the  
code does not match any of the codes determined to represent the set of referenced code  
lengths, as specified earlier in the compressed DHT (invalid DHT).

[0597]           •       A compressed format of a DHT (contents of a compressed data block  
15   with BTYPE equal 10 binary) is encountered which specifies a number of code lengths which  
is greater than the number of Huffman codes in the DHT, as specified by the sum of the  
values in the HLIT field, the HDIST field, and, e.g., 258. This is possible with improper uses  
of code lengths 16, 17, and 18, as examples (invalid DHT).

[0598]           •       A compressed format of a DHT (contents of a compressed data block  
20   with BTYPE equal 10 binary) is encountered which specifies a code length for the set of  
literal bytes, EOB symbol, and duplicate string lengths, which is less than the length required  
by the Huffman algorithm to specify a functional Huffman tree (invalid DHT).

[0599]           •       A compressed format of a DHT (contents of a compressed data block  
25   with BTYPE equal 10 binary) is encountered which specifies a code length for the set of  
duplicate string pointer distances, which is less than the length required by the Huffman  
algorithm to specify a functional Huffman tree (invalid DHT).

[0600]           •       A compressed data symbol, which is encountered in a compressed data  
block with BTYPE equal 10 binary, specifies a Huffman code which is not defined by the  
non-universal DHT derived from the compressed format of the DHT in the same block. In

this case, the number of bits of the second operand which are to be available to process, for the purpose of recognizing the general operand data exception, is model-dependent. More specifically, a model attempting to decode an undefined code may process, e.g., 15 bits prior to recognizing the exception, even though the exception could be determined after processing less bits.

[0601]           •       A compressed data symbol is encountered which is a duplicate string pointer and specifies a distance greater than the length of history available at the point of processing the symbol.

[0602]           •       A compressed data symbol, which is encountered in a compressed data block with BTYPE equal 01 binary, specifies an invalid code (e.g., a code of 11000110 or 11000111 binary for a duplicate string length, or a code of 11110 or 11111 binary for a duplicate string pointer distance). In this case, the number of bits of the second operand which are to be available to process, for the purpose of recognizing the general operand data exception, is model-dependent. More specifically a model attempting to decode an invalid code may process, e.g., 8 bits, in the case of a duplicate string length, or 5 bits, in the case of a duplicate string pointer distance, prior to recognizing the exception, even though the exception could be determined after processing less bits.

[0603]           When a general operand data exception is recognized, the operation is considered suppressed, even though operation ending supplemental code (OESC) 365 and model version number (MVN) fields 363 of the parameter block are updated to provide additional information associated with the exception.

[0604]           When a DFLTCC-CMPR or DFLTCC-XPND function is being executed and a general operand data exception is due to be recognized for the second operand, the result is that either the exception is recognized, or the operation ends with partial completion and condition code, e.g., 3 is set. If condition code 3 is set, the exception will be recognized when the instruction is executed again to continue processing the same operands and the exception condition still exists.

[0605]           Other Conditions include, for instance:

[0606] The execution of the instruction is interruptible. When an interruption occurs, the addresses in general registers  $R_1$  and  $R_2$ , the lengths in general registers  $R_1 + 1$  and  $R_2 + 1$ , and specific fields of the parameter block are updated, so that the instruction, when reexecuted, resumes at the point of interruption.

- 5 [0607] When a DFLTCC-CMPR or DFLTCC-XPND function is being executed and an access exception is due to be recognized for the first or second operand, the result is that either the exception is recognized, or the operation ends with partial completion and condition code, e.g., 3 is set. If condition code 3 is set, the exception will be recognized when the instruction is executed again to continue processing the same operands and the exception  
10 condition still exists.

[0608] As observed by this CPU, other CPUs, and channel programs, references to the parameter block, first, second, and third operands may be multiple-access references, accesses to these storage locations are not necessarily block concurrent, and the sequence of these accesses or references is undefined.

- 15 [0609] Results are unpredictable if the DFLTCC-CMPR or DFLTCC-XPND function is specified and any of the following apply, in one embodiment:

[0610] • The parameter block overlaps the first or second operand.

[0611] • The first operand overlaps the second operand.

- [0612] • The specified history buffer type (HBT) is circular and the third  
20 operand overlaps the first operand, the second operand, or the parameter block.

[0613] • The specified history buffer type (HBT) is in-line, the DFLTCC-CMPR function is specified, and the history overlaps the first operand or the parameter block.

[0614] • The specified history buffer type (HBT) is in-line, the DFLTCC-XPND function is specified, and the history overlaps the second operand or the parameter block.

- 25 [0615] In certain situations, despite ending the execution of the DEFLATE Conversion Call instruction with a CPU-determined number of bytes processed being zero,

data may have been stored to the first operand location, data may have been stored to the third operand location, when applicable, and corresponding change bits have been set, when applicable. In these cases, the contents of the parameter block and general registers have not been modified from original values. These situations may occur when the CPU performs a quiescing operation or a CPU retry while executing the DEFLATE Conversion Call instruction.

[0616] The following are example Resulting Condition Codes from executing the DEFLATE Conversion Call instruction:

[0617] 0 Normal completion

10 [0618] 1 The first operand length is insufficient to complete the operation

[0619] 2 The second operand length is insufficient to complete the operation (DFLTCC-XPND)

[0620] 3 CPU-determined amount of data processed

[0621] Program Exceptions:

15 [0622] • Access (fetch, operand 2, in-line history; fetch and store, parameter block, operand 1, operand 3)

[0623] • Data with DXC 0, general operand

[0624] • Operation (if the DEFLATE conversion facility is not installed)

[0625] • Specification

20 [0626] • Transaction constraint

[0627] Example priorities of execution for the DEFLATE CONVERSION CALL instruction are shown below:

[0628] 1.-6. Exceptions with the same priority as the priority of program-interruption conditions for the general case.

- [0629] 7.A Access exceptions for second instruction halfword.
- [0630] 7.B Operation exception.
- [0631] 7.C Transaction constraint.
- [0632] 8.A Specification exception due to invalid function code or invalid register  
5 number.
- [0633] 8.B Specification exception due to parameter block not designated on a 4  
K-byte boundary.
- [0634] 8.C Specification exception due to circular history buffer not designated on  
a 4 K-byte boundary.
- 10 [0635] 9. Access exceptions for an access to the parameter block.
- [0636] 10. General operand data exception when specified format of the parameter  
block is not supported by the mode.
- [0637] 11. Specification exception due to second operand length equal to zero and  
CF equal to zero at the beginning of the execution of the instruction.
- 15 [0638] 12. Condition code 1 due to first operand length equal to zero at the  
beginning of the execution of the instruction and DFLTCC-CMPR is specified.
- [0639] 13.A General operand data exception due to the history length field greater  
than 32,768 and the new task field is zero when DFLTCC-CMPR or DFLTCC-XPND is  
specified.
- 20 [0640] 13.B Access exceptions for an access to the first operand and the first  
operand length is non-zero.
- [0641] 13.C Access exceptions for an access to the second operand and the second  
operand length is non-zero.

[0642] 13.D Access exceptions for an access to in-line history specified at the beginning of the execution of the instruction.

[0643] 13.E Access exceptions for an access to the third operand.

[0644] 14.A General operand data exception due to conditions other than those  
5 included in items 10 and 13.A above.

[0645] 14.B Condition codes 1, 2 or 3 due to conditions other than those included in item 12 above.

[0646] 15. Condition code 0.

[0647] Prior to usage, the compressed format of a DHT is examined for the existence  
10 of general operand data exception conditions. When the length of the compressed format of a DHT is not precisely defined due to a general operand data exception condition, the interpreted length may depend on the condition, be model-dependent, and does not exceed, e.g., 286 bytes. As a result, when the DFLTCC-XPND function is specified and a compressed format of a DHT with a general operand data exception condition is encountered in the, e.g.,  
15 rightmost 286 bytes of the second operand, it is model-dependent whether the exception condition (priority 14.A) or condition code 2 (priority 14.B) is recognized.

[0648] Example programming notes are provided below:

[0649] 1. When compressing or uncompressing data, it may be more efficient overall when the operation is performed with a minimum number of times the DEFLATE  
20 Conversion Call instruction is executed. In other words, executing DFLTCC with a large operand may be more efficient than executing DFLTCC with small operands multiple times.

[0650] 2. For the compressing and uncompressing operations, when condition code 3 is set, the general registers used by the instruction and the parameter block have been updated such that the program can branch back to the instruction to continue the operation.

25 [0651] 3. In one embodiment, the DEFLATE Conversion Call instruction may be completed after performing a CPU-determined sub-portion of the processing specified by the

parameters of the instruction. When the instruction is completed after performing only a CPU-determined amount of processing instead of all specified processing, the instruction sets condition code 3. On such completion, the instruction address in the PSW (program status word) designates the next sequential instruction, and the operand parameters of the instruction have been adjusted so that the processing of the instruction can be resumed by branching back to the instruction to execute it again. When the instruction has performed all specified processing, it sets a condition code other than 3.

[0652] 4. When the DFLTCC-CMPR function is specified and the operation ends with a non-zero value in the sub-byte boundary (SBB) field of the parameter block, the operation included storing to the byte designated by the resulting first operand address. When the DFLTCC-XPND function is specified and the operation ends with a non-zero value in the SBB, the operation included fetching the byte designated by the resulting second operand address.

[0653] 5. When the operation ends with a non-zero condition code set, CSB field 392 of the parameter block may contain partially processed data, and it is expected that the program re-executes the instruction to resume the operation.

[0654] 6. Subsequent to an operation ending with a non-zero condition code set, and prior to re-executing the instruction for the purpose of resuming the operation, the program is not to modify any fields of the parameter block; otherwise results are unpredictable.

[0655] 7. When the DFLTCC-GDHT function is specified, the compressed representation of a DHT generated describes three proper-full Huffman code trees, according to the Huffman algorithm. That is, no under-full Huffman code trees are described. An under-full Huffman code tree is derived from a compressed representation of a DHT which specifies a code length for an element which is greater than the length required by the Huffman algorithm to specify a proper and functional Huffman tree.

[0656] When the DFLTCC-CMPR function is specified, HTT is one, and the compressed representation of the DHT includes a description of an under-full Huffman code

tree, the compressed data results can be transformed to the original uncompressed data by using the DFLTCC-XPND function, but not all decoders, which comply to the DEFLATE standard may be able to transform the results to the original uncompressed data. This may occur, for instance, when the compressed representation of a DHT specified by the program,  
5 for the DFLTCC-CMPR function was not generated as a result of performing the DFLTCC-GDHT function.

[0657] 8. When the DFLTCC-CMPR function ends with condition code 1 set, the result stored to sub-byte boundary (SBB) field 381 of the parameter block is 000 binary. Recognizing this scenario may be relevant to a program allocating output buffers for use with  
10 the DEFLATE Conversion Call instruction.

[0658] As described herein, in one aspect, a single instruction (e.g., a single architected machine instruction at the hardware/software interface, e.g., DEFLATE Conversion Call instruction) is provided to perform compress and/or decompress operations using a general-purpose processor. This instruction is, for instance, a hardware instruction  
15 defined in an Instruction Set Architecture (ISA). As a result, the complexity of the program related to the compressing and/or decompressing operations is reduced. Further, performance of the operations, and thus, the processor, is improved.

[0659] Advantageously, the DEFLATE Conversion Call instruction is dispatched, by, e.g., a programmer, on a general-purpose processor (e.g., a central processing unit, referred to  
20 herein as a processor), rather than a special-purpose processor, such as an I/O device, an application specific device connected through an I/O interface, or other types of special-purpose processors. Compared to a software implementation, executing the disclosed instruction requires significantly fewer execution cycles to perform the same operation. Further, compared to dispatching an operation to an I/O device, executing the disclosed  
25 instruction does not require I/O operations by an operating system and does not trigger the operating system to perform a task switch while waiting for the operation to complete.

[0660] Although various fields and registers are described, one or more aspects of the present invention may use other, additional or fewer fields or registers, or other sizes of fields and registers, etc. Many variations are possible. For instance, implied registers may be used

instead of explicitly specified registers or fields of the instruction and/or explicitly specified registers or fields may be used instead of implied registers or fields. Other variations are also possible.

[0661] One embodiment of using the DEFLATE Conversion Call instruction is  
5 described with reference to FIG. 17. In one example, a program executing on a processor, such as a general-purpose processor, specifies details of an operation to be performed in a parameter block in storage and specifies the location of the parameter block, STEP 1700. For instance, one or more of the fields of a parameter block (e.g., parameter block 340, 360 or 370) are provided or set, depending on the function to be performed. Further, the program  
10 specifies the operation to be performed (e.g., query, generate, compress, expand, etc.), STEP 1702. Additionally, the program specifies or updates the location and amount of input data in storage, STEP 1704, as well as the location and size of the result buffer in storage, STEP 1706.

[0662] Thereafter, the program executes the DEFLATE Conversion Call (DFLTCC)  
15 instruction, STEP 1708. In one example, the instruction is dispatched on a general-purpose processor. As examples, it is processed on the general-purpose processor, or, at least in part, is processed by hardware coupled to the general-purpose processor and accessible absent using an I/O interface.

[0663] Based on the instruction terminating, a determination is made as to whether the  
20 condition code resulting from execution is equal to a first defined value, e.g., 0, INQUIRY 1710. If the condition code is equal to the first defined value, then processing of the instruction is complete, STEP 1712. However, if the condition code is not equal to the first defined value, then a further determination is made as to whether the condition code is equal to a second defined value, e.g., 3, INQUIRY 1714. If the condition code is equal to the  
25 second defined value indicating there is additional data to be processed, then the instruction is re-executed, STEP 1708. However, if the condition code is not equal to the second defined value, then another determination is made as to whether the condition code is set to a third defined value, e.g., 1, INQUIRY 1716. If the condition code is set to the third defined value indicating the first operand length is insufficient, then processing continues with STEP 1706;

otherwise, the second operand length is insufficient for the function and processing continues with STEP 1704.

[0664] As indicated, the DEFLATE Conversion Call instruction may be executed multiple times to compress or decompress a single data stream. Therefore, in one aspect, the DEFLATE Conversion Call instruction includes an attribute which provides a mechanism for a program to declare a buffer (e.g., a 32 K-byte buffer), which is used to accumulate the history of uncompressed data processed during an operation which spans multiple executions of the DEFLATE Conversion Call instruction. The buffer is, for instance, a circular history buffer.

[0665] In one aspect, the DEFLATE Conversion Call instruction uses an indicator (e.g., a bit) in an implied register (e.g., GR0.56) to indicate use of a circular history buffer. When the circular history buffer is indicated and the specified function to be performed by the DEFLATE Conversion Call instruction is compressing or uncompressing data, a field of the instruction (e.g., R3) specifies the location in memory of, e.g., a 32 K-byte buffer, which the processor uses to fetch history from at the beginning of an operation and store history to at the end of an operation. The length of the history within the circular history buffer is specified by a field of a parameter block associated with the DEFLATE Conversion Call instruction (e.g., HL field 385), and the beginning of the history within the buffer is specified by an offset included in another field of the parameter block (e.g., HO field 386).

[0666] Further details of using a circular history buffer are described with reference to FIG. 18. In one example, a program executing on a processor, such as a general-purpose processor, specifies details of an operation to be performed in a parameter block in storage and specifies the location of the parameter block, STEP 1800. For instance, one or more of the fields of a parameter block (e.g., parameter block 360 or 370) are provided or set, depending on the function to be performed. Further, the program specifies the operation to be performed (e.g., compress, expand, etc.).

[0667] Further, in one example, the program allocates and specifies a location in memory of a pre-defined sized (e.g., 32 K-byte) circular buffer, STEP 1802. Additionally, the program places a portion of an uncompressed data stream into a buffer and specifies the

location and the size of the buffer as an input to the DEFLATE Conversion Call instruction, STEP 1804, and specifies or updates the location and size of a result buffer in storage, STEP 1806.

[0668] The DEFLATE Conversion Call instruction is then executed, STEP 1808.

5 Based on executing the instruction, the processor fetches history from, e.g., a circular history buffer, as an input to the operation, STEP 1820, and performs the specified operation, STEP 1822, as described herein. Further, the processor modifies the history in the circular history buffer as an output of the operation, STEP 1824. A determination is made as to whether the entire data stream has been processed, INQUIRY 1826. If not, then processing continues  
10 with STEP 1804. Otherwise, processing is complete.

[0669] Use of a circular history buffer provides the following, as examples:

[0670] When the size of the input or output buffer, specified for use with an individual execution of the DEFLATE Conversion Call instruction, is relatively small (e.g., 512 bytes), a history spanning multiple segments of buffered data, up to, e.g., 32 K-bytes, may be used as  
15 an input to the DEFLATE Conversion Call instruction, which processes a small number of bytes.

[0671] When the size of the input or output buffer, specified for use with an individual execution of the DEFLATE Conversion Call instruction, is relatively large (e.g., 128 K-bytes), a history of the prior segment of buffered data, up to, e.g., 32 K-bytes, may be used as  
20 a input to the DEFLATE Conversion Call instruction which is processing the first 32 K-bytes of data.

[0672] In both cases, more history is available to process data than would otherwise be available. As a result, the effectiveness of detecting duplicate strings is improved, resulting in improved overall compression ratios. This facilitates processing within the  
25 computing environment and improves performance.

[0673] One or more aspects of the present invention are inextricably tied to computer technology and facilitate processing within a computer, improving performance thereof. The use of a single architected machine instruction to perform compressing and/or decompressing

improves performance within the computing environment. The compressed/decompressed data may be used in many technical fields that manage and/or use data, such as in computer processing, medical processing, security, inventory control, etc. By providing optimizations in compressing/decompressing, these technical fields are improved by reducing execution  
5 time.

[0674] FIG. 19 depicts a system architecture 10100 in accordance with embodiments of the present invention. The system architecture 10100 includes a central processing unit (CPU) chip or integrated circuit (hereinafter referred to as a “CP”) 11000 that includes a backplane 11500, multiple processing units (PUs) 12000, an L3 cache and on-chip  
10 coherency unit (hereinafter referred to as an “on-chip coherency unit”) 13000, page buffer units (PBUs) 14000, an accelerator 15000 and a memory core (MC) 16000. The PUs 12000 include, for example, first PU 12000<sub>1</sub> and second PU 12000<sub>2</sub>. The PUs 12000, the on-chip coherency unit 13000, the PBUs 14000, the accelerator 15000 and the MC 16000 are supported and disposed on the backplane 11500. Each PU 12000 respectively communicates  
15 with the on-chip coherency unit 13000 by way of PU-L3 or first interfaces 13100. The on-chip coherency unit 13000 is communicative with external CPs and system controllers (SCs) via external busses 13200. Each PBU 14000 is respectively disposed in signal communication with peripheral component interconnect (PCI) interfaces 14100. The accelerator 15000 can be provided as an NXU (DEFLATE) accelerator. Each PBU 14000 and the accelerator 15000  
20 respectively communicates with the on-chip coherency unit 13000 by way of DMA-like or second interfaces 13300. The MC 16000 is communicative with external memory by way of memory interfaces 16100.

[0675] The accelerator 15000 is uniquely configured to execute certain functions and compression or decompression operations including, but not limited to DEFLATE  
25 operations and other lossless data compression or decompression algorithms. As noted above, DEFLATE, in particular, is an industry standard algorithm for compression of data that can be several gigabytes (GBs) in size where an application may only have small buffers at a time and the compression has to be completed in relatively small blocks that can be 1 megabyte (MB) or smaller.

[0676] While the DEFLATE operations and the other lossless data compression or decompression algorithms could be executed by any one of the PUs 12000, the accelerator 15000 is capable of completing the execution faster than any of the PUs 12000 even when taking into account the time required to communicate between the accelerator 15000 and other components of the CP 11000.

[0677] To the extent that the accelerator 15000 can be provided as an NXU (DEFLATE) accelerator, the accelerator 15000 is effectively a DMA-attached hardware accelerator that implements actual DEFLATE algorithm standards. When an application executes DEFLATE instructions, millicode operates the hardware of the accelerator 15000 on behalf of the application to perform the actual DEFLATE operations. As used herein, millicode is low level code that is transparent from an application perspective and is used to implement instructions or parts of instructions. For DEFLATE, millicode acquires and operates the accelerator 15000 on behalf of the application.

[0678] With reference to FIG. 20, the CP 11000, one or more additional or external CPs 20100 and one or more additional or external SCs 20200 can be provided in a drawer 21000. In such cases, the CP 11000, the one or more additional or external CPs 20100 and the one or more additional or external SCs 20200 can communicate with one another via external busses 20300 while the one or more additional or external SCs 20200 can also communicate with other features which are external with respect to the drawer 21000 via drawer interconnects 21100.

[0679] With reference to FIG. 21, the drawer 21000 can be provided as one of a plurality of drawers 30100 in which each individual drawer 21000 is communicative with each of the other individual drawers 21000.

[0680] With reference to FIG. 22, a method or a computer-implemented method of securely and efficiently executing data compression or decompression is provided for the system architecture 10100 described herein (the following description will relate to data compression for purposes of clarity and brevity). As shown in FIG. 22, the method is executable by the first and second processing units 12000<sub>1</sub> and 12000<sub>2</sub> acting in concert with one another and the accelerator 15000. The first processing unit 12000<sub>1</sub> determines whether to

operate in one of a synchronous mode or an asynchronous mode (40100) in order to conduct a data compression operation.

[0681] In an event the determining indicates that the asynchronous mode should be in effect, the first processing unit 12000<sub>1</sub> sets an ACF (40200). In this case, the data

5 compression operation is conducted in asynchronous mode. Here, the second processing unit 12000<sub>2</sub> issues a call to the accelerator 15000. In doing so, the second processing unit 12000<sub>2</sub> builds scatter gather lists (40300) to be fed to the accelerator 15000 to facilitate execution of the data compression by the accelerator 15000 in response to the setting of the ACF (40400). That is, DFLTCC millicode in the second processing unit 12000<sub>2</sub> takes absolute address from  
10 the scatter gather lists and feeds those absolute addresses to the accelerator 15000.

[0682] In an event the determining indicates that the synchronous mode should be in effect, the data compression operation is conducted in synchronous mode (40500). Here,

DFLTCC millicode in the first processing unit 12000<sub>1</sub> translates logical addresses of each page (in, e.g., 4k increments) of source and target buffers into absolute addresses and feeds  
15 those absolute addresses to the accelerator 15000.

[0683] With reference to FIG. 23, while the logical addresses of the source data and the logical addresses of the target data involved in the data compression or decompression operation are typically disposed consecutively in logical memory space, the absolute  
addresses of the source data and the absolute addresses of the target data are typically

20 disposed non-consecutively and, in some cases, non-sequentially in absolute memory space 50100. Thus, source address 1 with source length 1 occupies a first source region 51000 of absolute memory space 50100, source address 2 with source length 2 occupies a second source region 52000 of absolute memory space 50100 which is remote from the first source region 51000 and source address 3 with source length 3 occupies a third source region 53000  
25 of absolute memory space 50100 which is remote from the first and second source regions 51000 and 52000. Similarly, target address 1 with target length 1 occupies a first target region 53000 of absolute memory space 50100 and target address 2 with target length 2 occupies a second target region 54000 of absolute memory space 50100 which is remote from the first target region 53000.

[0684] With reference to FIG. 24, for the embodiment of FIG. 22, the scatter gather lists built by the second processing unit 12000<sub>2</sub> include a source scatter gather list 60100 and a target scatter gather list 60200. The source scatter gather list 60100 includes fields 61000 for each of the absolute addresses of the source data and fields 62000 for each of the lengths of the source data associated with each absolute address. The target scatter gather list 60200 includes fields 63000 for each of the absolute addresses of the target data and fields 64000 for each of the lengths of the target data associated with each absolute address.

[0685] The present invention may be a system, a method, and/or a computer program product at any possible technical detail level of integration. The computer program product may include a computer readable storage medium (or media) having computer readable program instructions thereon for causing a processor to carry out aspects of the present invention.

[0686] The computer readable storage medium can be a tangible device that can retain and store instructions for use by an instruction execution device. The computer readable storage medium may be, for example, but is not limited to, an electronic storage device, a magnetic storage device, an optical storage device, an electromagnetic storage device, a semiconductor storage device, or any suitable combination of the foregoing. A non-exhaustive list of more specific examples of the computer readable storage medium includes the following: a portable computer diskette, a hard disk, a random access memory (RAM), a read-only memory (ROM), an erasable programmable read-only memory (EPROM or Flash memory), a static random access memory (SRAM), a portable compact disc read-only memory (CD-ROM), a digital versatile disk (DVD), a memory stick, a floppy disk, a mechanically encoded device such as punch-cards or raised structures in a groove having instructions recorded thereon, and any suitable combination of the foregoing. A computer readable storage medium, as used herein, is not to be construed as being transitory signals per se, such as radio waves or other freely propagating electromagnetic waves, electromagnetic waves propagating through a waveguide or other transmission media (e.g., light pulses passing through a fiber-optic cable), or electrical signals transmitted through a wire.

[0687] Computer readable program instructions described herein can be downloaded to respective computing/processing devices from a computer readable storage

medium or to an external computer or external storage device via a network, for example, the Internet, a local area network, a wide area network and/or a wireless network. The network may comprise copper transmission cables, optical transmission fibers, wireless transmission, routers, firewalls, switches, gateway computers and/or edge servers. A network adapter card  
5 or network interface in each computing/processing device receives computer readable program instructions from the network and forwards the computer readable program instructions for storage in a computer readable storage medium within the respective computing/processing device.

[0688] Computer readable program instructions for carrying out operations of the  
10 present invention may be assembler instructions, instruction-set-architecture (ISA) instructions, machine instructions, machine dependent instructions, microcode, firmware instructions, state-setting data, configuration data for integrated circuitry, or either source code or object code written in any combination of one or more programming languages, including an object oriented programming language such as Smalltalk, C++, or the like, and  
15 procedural programming languages, such as the “C” programming language or similar programming languages. The computer readable program instructions may execute entirely on the user’s computer, partly on the user's computer, as a stand-alone software package, partly on the user's computer and partly on a remote computer or entirely on the remote computer or server. In the latter scenario, the remote computer may be connected to the user's  
20 computer through any type of network, including a local area network (LAN) or a wide area network (WAN), or the connection may be made to an external computer (for example, through the Internet using an Internet Service Provider). In some embodiments, electronic circuitry including, for example, programmable logic circuitry, field-programmable gate arrays (FPGA), or programmable logic arrays (PLA) may execute the computer readable  
25 program instruction by utilizing state information of the computer readable program instructions to personalize the electronic circuitry, in order to perform aspects of the present invention.

[0689] Aspects of the present invention are described herein with reference to flowchart illustrations and/or block diagrams of methods, apparatus (systems), and computer  
30 program products according to embodiments of the invention. It will be understood that each

block of the flowchart illustrations and/or block diagrams, and combinations of blocks in the flowchart illustrations and/or block diagrams, can be implemented by computer readable program instructions.

[0690] These computer readable program instructions may be provided to a processor of a general purpose computer, special purpose computer, or other programmable data processing apparatus to produce a machine, such that the instructions, which execute via the processor of the computer or other programmable data processing apparatus, create means for implementing the functions/acts specified in the flowchart and/or block diagram block or blocks. These computer readable program instructions may also be stored in a computer readable storage medium that can direct a computer, a programmable data processing apparatus, and/or other devices to function in a particular manner, such that the computer readable storage medium having instructions stored therein comprises an article of manufacture including instructions which implement aspects of the function/act specified in the flowchart and/or block diagram block or blocks.

[0691] The computer readable program instructions may also be loaded onto a computer, other programmable data processing apparatus, or other device to cause a series of operational steps to be performed on the computer, other programmable apparatus or other device to produce a computer implemented process, such that the instructions which execute on the computer, other programmable apparatus, or other device implement the functions/acts specified in the flowchart and/or block diagram block or blocks.

[0692] The flowchart and block diagrams in the Figures illustrate the architecture, functionality, and operation of possible implementations of systems, methods, and computer program products according to various embodiments of the present invention. In this regard, each block in the flowchart or block diagrams may represent a module, segment, or portion of instructions, which comprises one or more executable instructions for implementing the specified logical function(s). In some alternative implementations, the functions noted in the blocks may occur out of the order noted in the Figures. For example, two blocks shown in succession may, in fact, be executed substantially concurrently, or the blocks may sometimes be executed in the reverse order, depending upon the functionality involved. It will also be noted that each block of the block diagrams and/or flowchart illustration, and combinations of

blocks in the block diagrams and/or flowchart illustration, can be implemented by special purpose hardware-based systems that perform the specified functions or acts or carry out combinations of special purpose hardware and computer instructions.

[0693]           The descriptions of the various embodiments of the present invention have  
5   been presented for purposes of illustration, but are not intended to be exhaustive or limited to  
the embodiments disclosed. Many modifications and variations will be apparent to those of  
ordinary skill in the art without departing from the scope and spirit of the described  
embodiments. The terminology used herein was chosen to best explain the principles of the  
embodiments, the practical application or technical improvement over technologies found in  
10   the marketplace, or to enable others of ordinary skill in the art to understand the embodiments  
described herein.

## CLAIMS

1. A system architecture, comprising:

first and second processing units respectively communicative with an on-chip coherency unit; and

5 an accelerator communicative with the on-chip coherency unit,

the accelerator being configured to execute an operation responsive to a call issued by one of the first and second processing units,

the first processing unit being configured to set an asynchronous operation flag (AOF) to indicate that the second processing unit is to conduct an operation for the first processing  
10 unit,

the second processing unit being configured to respond to the AOF by building scatter gather lists and subsequently issuing the call and feeding the scatter gather lists to the accelerator to facilitate execution of the operation by the accelerator.

2. The system architecture according to claim 1, wherein the first and second  
15 processing units are each respectively provided as ones of multiple processing units respectively communicative with the on-chip coherency unit via first interfaces.

3. The system architecture according to claim 1, wherein the on-chip coherency unit is communicative with external processors and controllers.

4. The system architecture according to claim 1, further comprising page buffer  
20 units (PBUs) respectively communicative with the on-chip coherency unit via second interfaces and a memory core.

5. The system architecture according to claim 1, wherein the accelerator is communicative with the on-chip coherency unit via a second interface comprising a direct memory access (DMA)-like interface.

6. The system architecture according to claim 1, wherein the operation comprises a compression or decompression operation.

7. The system architecture according to claim 1, wherein the first processing unit is further configured to:

5 determine whether to operate in one of a synchronous mode or an asynchronous mode;

issue the call in an event the first processing unit determines that the synchronous mode should be in effect; and

set the AOF in an event the first processing unit determines that the asynchronous mode should be in effect.

10 8. The system architecture according to claim 1, wherein the scatter gather lists comprise absolute addresses and associated lengths for source data and target data of the operation.

9. The system architecture according to claim 8, wherein the absolute addresses of the source data are non-consecutive in absolute address space and the absolute addresses of  
15 the target data are non-consecutive in absolute address space.

10. The system architecture according to claim 8, wherein the absolute addresses of the source data are non-sequential and the absolute addresses of the target data are non-sequential.

11. A system architecture, comprising:

20 one or more drawers, each of which comprises one or more central processors (CPs) and one or more system controllers (SCs) and each of the one or more CPs comprising:

an accelerator configured to execute data compression or decompression faster than the first and second processing units responsive to a call issued by one of the first and second processing units,

the first processing unit being configured to set an asynchronous operation flag (AOF) to indicate that the second processing unit is to conduct a compression or decompression operation for the first processing unit,

the second processing unit being configured to respond to the AOF by building  
5 scatter gather lists and subsequently issuing the call and feeding the scatter gather lists to the accelerator to facilitate execution of the data compression or decompression by the accelerator.

12. The system architecture according to claim 11, wherein:

the one or more drawers are communitive with each other, and

10 the on-chip coherency unit of each of the one or more CPs of each of the one or more drawers is communicative with the one or more CPs and the one or more SCs of the corresponding drawer.

13. The system architecture according to claim 11, wherein the first and second processing units of each of the one or more CPs of each of the one or more drawers are  
15 provided as ones of multiple processing units respectively communicative with a corresponding on-chip coherency unit via first interfaces.

14. The system architecture according to claim 11, wherein each of the one or more CPs of each of the one or more drawers further comprises page buffer units (PBUs) respectively communicative with a corresponding on-chip coherency unit via second  
20 interfaces and a memory core.

15. The system architecture according to claim 11, wherein the accelerator of each of the one or more CPs of each of the one or more drawers is communicative with a corresponding on-chip coherency unit via a second interface comprising a direct memory access (DMA)-like interface.

25 16. The system architecture according to claim 11, wherein the compression or decompression operation comprises a DEFLATE operation.

17. The system architecture according to claim 11, wherein the first processing unit is further configured to:

determine whether to operate in one of a synchronous mode or an asynchronous mode;

5 issue the call in an event the first processing unit determines that the synchronous mode should be in effect; and

set the ACF in an event the first processing unit determines that the asynchronous mode should be in effect.

18. The system architecture according to claim 11, wherein the scatter gather lists comprise absolute addresses and associated lengths for source data and target data of the data  
10 compression or decompression.

19. The system architecture according to claim 18, wherein the absolute addresses of the source data are non-consecutive and the absolute addresses of the target data are non-consecutive.

20. The system architecture according to claim 19, wherein the absolute addresses  
15 of the source data are non-sequential in absolute address space and the absolute addresses of the target data are non-sequential in absolute address space.

21. A computer program product, comprising:

a processing circuit; and

20 a memory unit having executable instructions stored thereon, which, when executed cause the processing circuit to execute a method of securely and efficiently executing data compression or decompression, the method comprising:

determining whether to operate in one of a synchronous mode or an asynchronous mode;

25 setting an asynchronous operation flag (AOF) in an event the determining indicates that the asynchronous mode should be in effect; and

building scatter gather lists to be fed to an accelerator to facilitate execution of the data compression or decompression by the accelerator in response to the setting of the AOF.

22. The computer program product according to claim 21, wherein:

5 the scatter gather lists comprise absolute addresses and associated lengths for source data and target data of the data compression or decompression, and

the absolute addresses of the source data are non-consecutive and the absolute addresses of the target data are non-consecutive.

23. A computer-implemented method of securely and efficiently executing data  
10 compression or decompression, the computer-implemented method comprising:

determining whether to operate in one of a synchronous mode or an asynchronous mode;

setting an asynchronous operation flag (AOF) in an event the determining indicates that the asynchronous mode should be in effect; and

15 building scatter gather lists to be fed to an accelerator to facilitate execution of the data compression or decompression by the accelerator in response to the setting of the AOF.

24. The computer-implemented method according to claim 21, wherein:

the scatter gather lists comprise absolute addresses and associated lengths for source data and target data of the data compression or decompression, and

20 the absolute addresses of the source data are non-consecutive and the absolute addresses of the target data are non-consecutive.

25. A system architecture, comprising:

first and second processing units respectively communicative with an on-chip coherency unit; and

an accelerator communicative with the on-chip coherency unit,

the accelerator being configured to execute data compression or decompression faster than the first and second processing units responsive to a call issued by one of the first and second processing units,

5           the first processing unit being configured to execute a method comprising:

                  determining whether to operate in one of a synchronous mode or an asynchronous mode; and

                  setting an asynchronous operation flag (AOF) in an event the determining indicates that the asynchronous mode should be in effect, and

10           the second processing unit being configured to execute a method comprising building scatter gather lists to be fed to the accelerator to facilitate execution of the data compression or decompression by the accelerator in response to the setting of the AOF.




FIG. 1A




FIG. 1B




FIG. 2




FIG. 3A




FIG. 3B

|     | CODE | FUNCTION      | PARAMETER<br>BLOCK SIZE (BYTES) |
|-----|------|---------------|---------------------------------|
| 313 | 0    | DFLTCC - QAF  | 32                              |
| 315 | 1    | DFLTCC - GDHT | 384                             |
| 317 | 2    | DFLTCC - CMPR | 1536                            |
| 319 | 4    | DFLTCC - XPND | 1536                            |

FIG. 3C




FIG. 3D




FIG. 3E




FIG. 3F




FIG. 3G




FIG. 3H




FIG. 3I




FIG. 3J






FIG. 3K



370

PARAMETER BLOCK FOR DFLTCC - CMPR + DFLTCC - XPND

| PARAMETER BLOCK<br>VERSION NUMBER <u>362</u>          |                                   |                                      |                                   | MODEL<br>VERSION NUMBER <u>363</u>     |                                     |                     | CONTINUATION<br>FLAG <u>373</u>    |                                            |                                        |                                        |
|-------------------------------------------------------|-----------------------------------|--------------------------------------|-----------------------------------|----------------------------------------|-------------------------------------|---------------------|------------------------------------|--------------------------------------------|----------------------------------------|----------------------------------------|
| NEW<br>TASK<br><u>374</u>                             | CHECK<br>VALUE<br>TYPE <u>375</u> | HUFFMAN<br>TABLE<br>TYPE <u>376</u>  | BLOCK<br>CONT.<br>FLAG <u>377</u> | BLOCK<br>CLOSING<br>CONTROL <u>378</u> | BLOCK<br>HEADER<br>FINAL <u>379</u> | DHTGC<br><u>364</u> | SUB-BYTE<br>BOUNDARY<br><u>381</u> | OPER<br>ENDING<br>SUPP. CODE<br><u>365</u> | INCOMP<br>FUNC<br>STATUS<br><u>383</u> | INCOMP<br>FUNC<br>LENGTH<br><u>384</u> |
| HISTORY LENGTH<br><u>385</u>                          |                                   |                                      |                                   | HISTORY OFFSET<br><u>386</u>           |                                     |                     |                                    |                                            |                                        |                                        |
| CHECK VALUE<br><u>387</u>                             |                                   | END-OF-BLOCK<br>SYMBOL<br><u>388</u> |                                   | END-OF-BLOCK<br>LENGTH<br><u>389</u>   |                                     |                     |                                    |                                            |                                        |                                        |
|                                                       |                                   |                                      |                                   |                                        |                                     |                     |                                    |                                            |                                        |                                        |
|                                                       |                                   |                                      |                                   |                                        |                                     |                     |                                    |                                            |                                        |                                        |
|                                                       |                                   |                                      |                                   |                                        |                                     |                     |                                    |                                            |                                        |                                        |
| COMPRESSED DYNAMIC-HUFFMAN TABLE LENGTH<br><u>366</u> |                                   |                                      |                                   |                                        |                                     |                     |                                    |                                            |                                        |                                        |
| COMPRESSED DYNAMIC-HUFFMAN TABLE<br><u>367</u>        |                                   |                                      |                                   |                                        |                                     |                     |                                    |                                            |                                        |                                        |
| CONTINUATION STATE BUFFER<br><u>392</u>               |                                   |                                      |                                   |                                        |                                     |                     |                                    |                                            |                                        |                                        |

FIG. 3L








FIG. 6




FIG. 7




FIG. 8




FIG. 9

1000

|                     |               |                    |                                                                                               |
|---------------------|---------------|--------------------|-----------------------------------------------------------------------------------------------|
| 1002<br>}<br>IABLK1 | L             | 1,PBLKADDR         | Load parameter block address into GR1                                                         |
|                     | ILL           | 0,2                | Set GR0(56:63)=02 for DFLTCC-CMPR function                                                    |
|                     | L             | 2,CDSBADDR         | Load compressed data set output address into GR2                                              |
|                     | L             | 4,UD1ADDR          | Load uncompressed data input address into GR4                                                 |
|                     | DFLTCC<br>BRC | 2,4,10<br>1,IABLK1 | Compress data into block1<br>If operation ends with CC=3, branch back to resume the operation |
| 1004<br>}<br>IABLK2 | :             |                    |                                                                                               |
|                     | :             |                    |                                                                                               |
|                     | L             | 6,UD2ADDR          | Load next uncompressed data input address into GR6                                            |
|                     | DFLTCC        | 2,6,10             | Compress data into block2                                                                     |
|                     | BRC           | 1,IABLK2           | If operation ends with CC=3, branch back to resume the operation                              |
| IABLK3              | :             |                    |                                                                                               |
|                     | :             |                    |                                                                                               |
|                     | L             | 8,UD3ADDR          | Load next uncompressed data input address into GR8                                            |
|                     | DFLTCC        | 2,8,10             | Compress data into block3                                                                     |
|                     | BRC           | 1,IABLK3           | If operation ends with CC=3, branch back to resume the operation                              |
|                     | :             |                    |                                                                                               |
|                     | :             |                    |                                                                                               |
|                     | :             |                    |                                                                                               |
|                     | :             |                    |                                                                                               |
|                     | :             |                    |                                                                                               |

Sample of a Program Compressing Data into Three Blocks of a Compressed Data Set

FIG. 10

|                                 |                                              |                                                                                                                                     |                                                                                                        |
|---------------------------------|----------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------|
| parameter block field<br>1100 } | field value at start of operation ~1102      | field value at end of ~1104 operation when condition code 1, 2, or 3 is set (also the value at the start of resuming the operation) | field value at end of operation when condition code 0 is set ~1106                                     |
| new task (NT)                   | 1 (specified by the program)                 | 0                                                                                                                                   | 0                                                                                                      |
| continuation flag (CF)          | initialized to zero by the program           | 1                                                                                                                                   | 0                                                                                                      |
| continuation state buffer (CSB) | initialized to all zeros by the program      | internal-state data recorded at the end of the operation, which is required when the operation later resumes                        | undefined                                                                                              |
| check value                     | treated as an initial value, since NT is one | value generated from portion of data processed since the start of the operation and starting value for resuming the operation       | value generated from all data processed by the operation and starting value for a subsequent operation |
| sub-byte boundary (SBB)         | treated as equal to zero, since NT is one    | applies to the last byte processed by the operation and the first byte when the operation later resumes                             | applies to the last byte processed by the operation and the first byte for a subsequent operation      |

Parameter Block Contents for a DFL TCC-CMPR Function Operating on the First Compressed Data Block of a Set

FIG. 11

| parameter block field           | field value at start of operation                                                           | field value at end of operation when condition code 1, 2, or 3 is set (also the value at the start of resuming the operation)                 | field value at end of operation when condition code 0 is set                                           |
|---------------------------------|---------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------|
| new task (NT)                   | 0 (specified by the program)                                                                | 0                                                                                                                                             | 0                                                                                                      |
| continuation flag (CF)          | initialized to zero by the program                                                          | 1                                                                                                                                             | 0                                                                                                      |
| continuation state buffer (CSB) | initialized to all zeros by the program                                                     | internal-state data recorded at the end of the operation, which is required when the operation later resumes                                  | undefined                                                                                              |
| check value                     | value generated from prior operation (preserved by the program)                             | value generated from portion of data processed since the start of the current operation and starting value for resuming the current operation | value generated from all data processed by the operation and starting value for a subsequent operation |
| sub-byte boundary (SBB)         | value associated with last byte processed by the prior operation (preserved by the program) | applies to the last byte processed by the current operation and the first byte when the current operation later resumes                       | applies to the last byte processed by the operation and the first byte for a subsequent operation      |

*Parameter Block Contents for a DFL TCC-CMPR Function Operating on the Second Compressed Data Block of a Set*

**FIG. 12**

1300

|       |            |                                                                   |
|-------|------------|-------------------------------------------------------------------|
| L     | 1,PBLKADDR | Load parameter block address into GR1                             |
| IILL  | 0,4        | Set GR0(56:63)=04 for DFLTCC-XPND function                        |
| L     | 2,UDADDR   | Load uncompressed data output address into GR2                    |
| L     | 4,CDSBADDR | Load uncompressed data set input address into GR4                 |
| IASET | 2,4,10     | Uncompress all blocks of the data set                             |
|       | 1,IASET    | If operation ended with CC=3, branch back to resume the operation |
|       | DFLTCC     |                                                                   |
|       | BRC        |                                                                   |
|       | :          |                                                                   |
|       | :          |                                                                   |

Sample of a Program Uncompressing Data from a Compressed Data Set

FIG. 13




FIG. 14A





FIG. 14B



**Explanation:**

shaded areas: bytes of history

SB source-begin address: location specified by  $R_2$

SE source-end address: location specified by  $R_2 + (R_2 + 1) - 1$

HB history-begin address: location specified by  $R_2 - HL$

HE history-end address: location specified by  $R_2 - 1$

LR least recently processed byte in history

MR most recently processed byte in history

BP bytes of uncompressed data processed

HL history length

*In-line History before and after executing DFLTCC-CMPR multiple times*

FIG. 14C




FIG. 15A




FIG. 15B




FIG. 15C




FIG. 15D




FIG. 15E




FIG. 16A





FIG. 16B



**Explanation:**

- shaded areas: bytes of history
- TB target-begin address: location specified by  $R_1$
- TE target-end address: location specified by  $R_1 + (R_1 + 1) - 1$
- HB history-begin address: location specified by  $R_1 - HL$
- HE history-end address: location specified by  $R_1 - 1$
- LR least recently processed byte in history
- MR most recently processed byte in history
- BP bytes of uncompressed data processed
- HL history length
- In-line History before and after executing DFLTCC-XPND multiple times*

FIG. 16C

19/24




FIG. 17

20/24




FIG. 18




FIG. 19




FIG. 21




FIG. 20




FIG. 22




FIG. 23

FIG. 24

## INTERNATIONAL SEARCH REPORT

International application No  
PCT/EP2020/053383

## A. CLASSIFICATION OF SUBJECT MATTER

INV. G06F9/38

ADD. G06F9/30 H03M7/30

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)

H03M G06F

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

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

EPO-Internal, WPI Data

## C. DOCUMENTS CONSIDERED TO BE RELEVANT

| Category* | Citation of document, with indication, where appropriate, of the relevant passages                                                                            | Relevant to claim No. |
|-----------|---------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------|
| A         | US 2017/187388 A1 (SATPATHY SUDHIR K [US] ET AL) 29 June 2017 (2017-06-29)<br>figures 1,2<br>paragraph [0032] - paragraph [0037]<br>paragraph [0046]<br>----- | 1-25                  |
| A         | US 2007/109153 A1 (MA KEVIN [US] ET AL) 17 May 2007 (2007-05-17)<br>the whole document<br>-----                                                               | 1-25                  |
| A         | WO 99/34273 A2 (LSI LOGIC CORP [US]) 8 July 1999 (1999-07-08)<br>page 1 - page 3<br>-----                                                                     | 1-25                  |
| A         | US 7 971 030 B2 (MMAGIX TECHNOLOGY LTD [GB]) 28 June 2011 (2011-06-28)<br>the whole document<br>-----<br>-/-                                                  | 1-25                  |



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

"&amp;" document member of the same patent family

Date of the actual completion of the international search

8 May 2020

Date of mailing of the international search report

19/05/2020

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

De Poy, Iker

## INTERNATIONAL SEARCH REPORT

International application No

PCT/EP2020/053383

| C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT |                                                                                                                                        |                       |
|------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------|-----------------------|
| Category*                                            | Citation of document, with indication, where appropriate, of the relevant passages                                                     | Relevant to claim No. |
| A                                                    | US 2016/233882 A1 (AGARWAL KANAK B [US] ET AL) 11 August 2016 (2016-08-11)<br>figure 1<br>paragraph [0030] - paragraph [0035]<br>----- | 1-25                  |

# INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/EP2020/053383

| Patent document<br>cited in search report | Publication<br>date | Patent family<br>member(s) | Publication<br>date           |
|-------------------------------------------|---------------------|----------------------------|-------------------------------|
| US 2017187388                             | A1                  | 29-06-2017                 | CN 108292222 A 17-07-2018     |
|                                           |                     |                            | DE 112016006059 T5 06-09-2018 |
|                                           |                     |                            | TW 201732567 A 16-09-2017     |
|                                           |                     |                            | US 2017187388 A1 29-06-2017   |
|                                           |                     |                            | WO 2017112351 A1 29-06-2017   |
| -----                                     |                     |                            |                               |
| US 2007109153                             | A1                  | 17-05-2007                 | NONE                          |
| -----                                     |                     |                            |                               |
| WO 9934273                                | A2                  | 08-07-1999                 | NONE                          |
| -----                                     |                     |                            |                               |
| US 7971030                                | B2                  | 28-06-2011                 | AU 2003250575 A1 25-02-2004   |
|                                           |                     |                            | EP 1543418 A1 22-06-2005      |
|                                           |                     |                            | US 2005235134 A1 20-10-2005   |
|                                           |                     |                            | US 2011289276 A1 24-11-2011   |
|                                           |                     |                            | US 2013311723 A1 21-11-2013   |
|                                           |                     |                            | US 2016210155 A1 21-07-2016   |
|                                           |                     |                            | WO 2004015572 A1 19-02-2004   |
| -----                                     |                     |                            |                               |
| US 2016233882                             | A1                  | 11-08-2016                 | US 2015066878 A1 05-03-2015   |
|                                           |                     |                            | US 2016233882 A1 11-08-2016   |
| -----                                     |                     |                            |                               |