



- (51) **International Patent Classification:**  
*H03M 7/30* (2006.01) *H04L 29/08* (2006.01)
- (21) **International Application Number:**  
PCT/IB201 1/055441
- (22) **International Filing Date:**  
2 December 2011 (02.12.2011)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (71) **Applicant (for all designated States except US):** CANON KABUSHIKI KAISHA [JP/JP]; 3-30-2, Shimomarucho, 3-chome, Ohta-ku, Tokyo (JP).
- (72) **Inventors; and**
- (75) **Inventors/Applicants (for US only):** FABLET, Youenn [FR/FR]; La Roulais, F-35390 La Dominelais (FR). BEL-LESSORT, Romain [FR/FR]; 5 rue Jean Brulelou dit Brito, F-35000 Rennes (FR). RUELLAN, Herve [FR/FR]; 49 rue Duhamel, F-35000 Rennes (FR).
- (74) **Agent:** LE TOURNEAU, Augustin; Santarelli, BP 237, 14 avenue de la Grande-Armee, F-75822 Paris Cedex 17 (FR).
- (81) **Designated States (unless otherwise indicated, for every kind of national protection available):** AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.
- (84) **Designated States (unless otherwise indicated, for every kind of regional protection available):** ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK,

[Continued on nextpage]

- (54)
- Title:**
- METHODS AND DEVICES FOR ENCODING AND DECODING MESSAGES




Fig. 3

- (57) **Abstract:** Methods and devices for encoding or decoding messages, each message including a list of information items. The encoding method comprises determining a first list of indexes associated with the information items that are already indexed in a local indexing table and a second list of literal values of the other information items of the list not yet indexed in said local indexing table; encoding the indexes of the first list; binary compressing at least a serialized binary representation of the literal values of the second list; and concatenating the first list and the second list together to obtain an encoded bitstream of the information items. When the messages are sent over a plurality of connections, a global table is shared between the connections to store the indexed items of information; and a local indexing table for each connection associates indexes with references to an entry of the shared global table.

SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).  
**Published:**  
— -with international search report (Art. 21(3))  
**Declarations under Rule 4.17:**  
— of inventorship (Rule 4.17(iv))

## METHODS AND DEVICES FOR ENCODING AND DECODING MESSAGES

5           The present invention concerns methods of encoding and of decoding messages, as well as corresponding encoding and decoding devices.

          The invention belongs to the field of network communication, and in particular to the field of data compression used when sending messages over a communications network.

10           In messages to be sent there are often lists or groups of items of information that are compressed at the encoder and decompressed at the decoder. This is for example the case for HTTP (standing for Hypertext Transfer Protocol) where HTTP payload is compressed, as well as for SPDY protocol where HTTP headers are compressed.

15           SPDY is a protocol that intends to improve the well-known HTTP. The SPDY protocol is referred to below to illustrate the present invention, although the invention may apply to other configurations. In particular, it may apply to a large number of data transmission protocols where a list of information items is compressed, such as the above cited HTTP and SPDY protocols.

20           HTTP is commonly used to request and send web pages, and is based on a client/server architecture, wherein the client sends requests, namely HTTP requests, to the server, and the server replies to the client's requests with responses, namely HTTP responses.

          Requests and responses are messages that comprise various parts,  
25   among which are non-compressed HTTP headers and compressed HTTP payload.

          Figure 1 represents an example of six HTTP headers that are successively listed in an HTTP request, and then in an HTTP response.

          An HTTP header consists of a name (in bold in the Figure) along with a corresponding value.

30           For instance, in "Host: en.wikipedia.org", Host is the header name, and its value is "en.wikipedia.org". This header is used to indicate the host of the requested resource (for instance, the Wikipedia page describing HTTP, available at <http://en.wikipedia.org/wiki/HTTP>). HTTP headers are well-known to one skilled in the art, and therefore are not further detailed here.

In the first versions of HTTP, a TCP/IP connection was established for each HTTP request/response exchange.

SPDY is a protocol that improves HTTP in several ways.

First, it enables several HTTP requests and responses to be sent over a unique TCP/IP connection, thus defining a set of HTTP transmissions therein. In this way, all the components of a web page (HTML documents, images, JavaScript, etc.) may share the same TCP/IP connection, thus speeding up the web page loading.

Secondly, SPDY implements compression of the HTTP headers exchanged over the shared TCP/IP connection, using the Zlib Deflate algorithm (also known through the "zip" format). This binary compression reduces the network load.

The Deflate algorithm performs compression of a serialized binary representation of the HTTP headers, by searching (using a sliding window) and replacing duplicate strings of the binary representation with back references thereto, and replacing symbols of the back references with Huffman codes.

A serialized binary representation of the HTTP headers results from the serialization of the HTTP headers as a stream of bits (or bytes).

Compression gains obtained by SPDY are acceptable.

However, speed performance is observed to be poor, while memory consumption may become disadvantageously high.

The poor speed performance is due to the potentially large size of the serialized binary representation, resulting from a large number of headers.

High memory consumption may arise in the case of a HTTP server handling simultaneously a plurality of connections towards a plurality of clients. Independently encoding the headers of the messages transmitted over the different connections requires a large amount of memory, in particular when the number of connections increases: one encoder or compressor is assigned to each connection and a dedicated memory area is reserved for each encoder or compressor.

In this context, there is a need for a new approach for compressing a list of information items (such as HTTP headers) that achieves similar size reduction while substantially reducing the memory and processing costs.

The present invention has been devised to address one or more of the foregoing concerns.

According to a first aspect of the invention there is provided a method of encoding messages, a message including an initial list of items of information, the method comprising:

determining a first group of indexes associated with the information items of the initial list that are already indexed in a local indexing table and a second group of literal values of the other information items of the initial list not yet indexed in said local indexing table;

- 5                    encoding the indexes of the first group;  
                     binary compressing at least a serialized binary representation of the literal values of the second group; and  
                     concatenating the first group and the second group together to obtain an encoded bitstream of the information items.

- 10                   The invention substantially improves the processing speed, for example up to twice as fast as SPDY when the HTTP headers are highly redundant. It may also improve the memory consumption.

- This is achieved by indexing items prior to conventional binary compression and grouping based on the indexing step. Indexing allows very fast redundancy search  
15                   with a limited memory usage at the encoder, while at the decoder, the decoding is rendered more efficient with faster building of the information items.

                     This is because the indexing technique is based on a highly focused redundancy search while binary compression such as Deflate uses a brute-force exhaustive redundancy search which is slower.

- 20                   According to a second aspect of the invention there is provided a method of decoding encoded messages, the method comprising:

- obtaining encoded data from a bitstream, the encoded data comprising a first part concatenated with a second part;  
                     decoding the first part to obtain a first group of indexes associated with  
25                   information items already indexed in a local indexing table;  
                     binary decompressing the second part to obtain a serialized binary representation of a second group of literal values corresponding to information items not yet indexed in said local indexing table; and  
                     providing a decoded list made of the information items indexed by the  
30                   indexes of the first group and of the information items literally decoded from the second group.

                     This method provides advantages similar to the above method, but at the decoder. It is to be noted that the information items of the first and second groups are of the same time, for example HTTP headers.

According to a third aspect of the invention there is provided a method of encoding messages to send over a plurality of connections, a message including a list of items of information, the method comprising:

5 encoding the list of information items of a message using a local indexing table specific to the connection to which the message belongs;

wherein a global table is shared between the connections and comprises items of information; and the local indexing table of each connection associates indexes with references to an entry of the shared global table.

10 This aspect of the invention optimizes memory consumption. This is because the global table which stores the majority of data (here the information items) is shared between all the connections enabling each information item to be stored only once, while the indexing is specific to each connection providing the shortest indexes for each connection. Thus the global table makes it possible to share the memory cost of the indexing tables between the connections.

15 Conversely, according to a fourth aspect of the invention there is provided a method of decoding encoded messages received over a plurality of connections, an encoded message including an encoded list of items of information, the method comprising:

20 decoding the encoded list of information items of an encoded message using a local indexing table specific to the connection to which the encoded message belongs;

wherein a global table is shared by all the connections and comprises items of information; and the local indexing table of each connection associates indexes with references to an entry of the shared global table.

25 This method provides advantages similar to the above coding method, but at the decoder.

Optional features of the invention are further defined in the dependent appended claims. Below, when discussing certain features with respect to the encoding or compressing, one skilled in the art will easily understand that the corresponding  
30 features at the encoder may also be used with respect to decoding or decompressing in the present invention.

In one embodiment of the invention, the first group of indexes is prior to the second group of literal values in the encoded bitstream. At the decoder, this means that the first part encoding indexes is prior to the second part (of literal values) of the  
35 bitstream.

This configuration gives access to a large number of information items with few encoded data to decode by the decoder, for example to access HTTP headers in SPDY. This is because the majority of data representing each indexed information item is stored in the local indexing table and not encoded in the bitstream.

5 In another embodiment, the step of concatenating generates a serialized binary representation of the concatenated first and second groups, and the step of binary compressing compresses the whole serialized binary representation. A shorter sliding window for binary compression may then be used since the redundancies for the indexes occur more often in the serialized binary representation. A shorter sliding  
10 window means memory saving (fewer bits to buffer) as well as processing speed improvement (fewer bits to search through).

In yet another embodiment, the items of information are HTTP headers of an HTTP message. This is the case for SPDY for example.

In particular, an HTTP header is made of at least an HTTP name and an  
15 HTTP value, and an indexed HTTP header of the first group comprises an index indexing the HTTP name and comprises the literal HTTP value. This takes advantage of the high redundancy between the HTTP header names (which are generally used in consecutive HTTP messages) while their values may vary highly from one message to another. The same approach may be used for any information item made of a name  
20 and a value.

According to a particular feature, an HTTP header is made of at least an HTTP name and an HTTP value, and the method further comprises associating an indexing level flag with an indexed HTTP header of the first group, said indexing level flag taking, depending on the indexed HTTP header, either a first full indexing state  
25 indicating that the indexed HTTP header has the same HTTP name and HTTP value as the corresponding information item in the local indexing table, or a second partial indexing state indicating that the indexed HTTP header has the same HTTP name as the corresponding information item in the local indexing table but a different HTTP value. This makes it possible to take advantage of the high redundancy of some slowly-  
30 varying HTTP headers. For example, it may be advantageous to index a time HTTP header updated each second, since several messages may be sent within a single second. Then for the messages in the following second, the partial indexing state of the flag is used.

In another embodiment of the invention, the method further comprises  
35 updating the local indexing table by associating, within that local indexing table, at least

one information item of the second group with a new index. This makes it possible to feed the local indexing table in such that a new occurrence of the same information item can be indexed.

5 In particular, the method may further comprise associating, in the encoded stream, an indexing flag with each binary-compressed information item of the second group that is associated with a new index during said updating. This makes it possible to indicate to the decoder an indexing strategy that it may not know.

10 According to another particular feature, the method may further comprise determining a frequency of occurrence of an item of information of the second group, and deciding to index said item of information in the local indexing table depending on the determined frequency. Thanks to this provision, only the most often occurring items are indexed which leads to an optimized trade-off between memory used to store the local indexing table and processing speed improvement due to the use of indexes.

15 In yet another embodiment of the invention, the local indexing table associates each index with a reference to an entry of a global table, said entry of the global table comprising the item of information associated with said index. This makes it possible to share the same global table for several connections over which messages are transmitted. In this context, a local indexing table will be used for each connection.

20 According to a particular feature, the local indexing table comprises at least one item of information indexed by reference to the global table and at least one item of information indexed within the local indexing table alone. A refined indexing strategy may then be implemented that distinguishes certain information items shared by several connections or that are highly frequent (which can be stored by reference to the shared global table) from certain information items specific to a particular connection  
25 (which may then be stored within the local indexing table specific to that connection).

In particular, the method may further comprise determining a frequency of occurrence of an item of information, and deciding to index said item of information in the local indexing table alone or by reference to the global table depending on the determined frequency.

30 According to another particular feature, the global table is shared between a plurality of compressors (corresponding to a plurality of connections) and comprises items of information from messages processed by said plurality of compressors, each compressor having its local indexing table to index, by reference to the shared global table, items of information from messages that it processes. This provision explicitly

shows that the global table is shared between the compressors, while the local indexing tables are specific to each compressor.

According to yet another particular feature, determining the first and second groups comprises searching for entries in the global table that correspond to the information items of the initial list and searching, in the local indexing table, for references to the entries found. Such searching is adapted to the use of the shared global table.

According to another embodiment of the invention, the method further comprises choosing between a first encoding of the initial list of information items using said steps of determining, concatenating and binary compressing, and a second encoding of the initial list of information items using only binary compression of a serialized binary representation of the item literal values (i.e. activating or not activating the indexing);

a sliding binary window used for the binary compression of the first encoding being of lower size than a sliding binary window used for the binary compression of the second encoding.

Thanks to this configuration memory use is reduced when using the indexing of information items. This is because thanks to the indexing the redundancies within the serialized binary representation occur more often, such that a shorter sliding window is sufficient to detect those redundancies.

According to yet another embodiment of the invention, the global table is shared between a plurality of network connections and referred to by a corresponding plurality of respective local indexing tables, and

the method further comprises, at some point during the encoding, migrating current connections from said shared global table to a new (shared) global table.

This provision of migrating the global table provides saving of memory. This is because, as the messages are exchanged and information items are indexed, the number of indexed information items increases in the shared global table. The migration makes it possible to purge unused indexes from that table, thus saving memory.

That means that the features of the above third aspect of the invention and its optional features may be implemented in the above-defined first aspect of the invention. Conversely, the encoding process defined above in the first aspect of the invention and its optional features may be implemented by some or all the plurality of network connections.

According to a particular feature, the new global table may be created upon receiving a request for a new connection, depending on at least one criterion relating to the currently shared global table (for example the age of that table). It is then decided, as of making this new connection, to work on a new clean global table for that new connection. The new clean global table is a purged version of the previous shared global table.

In particular, the method may further comprise removing entries from the currently shared global table based on the amount of memory used to store both currently shared and new global tables. This makes it possible to comply with memory constraints when using both the new and current global tables before the current global table is abandoned and deleted.

According to another particular feature, migrating current connections is triggered upon receiving a request to end a current connection.

In particular, migrating current connections comprises removing entries used by the current connections from the currently shared global table; adding those removed entries to the new global table; and then deleting said currently shared global table from memory. This provision ensures that the existing indexes are kept for the current connections while freeing memory space.

According to yet another embodiment of the invention, the global table is shared between a plurality of network connections and referred to by a corresponding plurality of respective local indexing tables, each indexed information item in the shared global table being associated with a connection counter that counts the number of connections using that indexed information item, and

the method further comprises performing a memory saving operation based on connection counters of the indexed information items.

The use of the connection counter makes it possible to easily detect when an entry (i.e. an indexed information item) of the shared global table is no longer used by any connection. In this situation this entry may be deleted from the table, thus freeing memory space.

In particular, the memory saving operation comprises removing the indexed information item from the shared global table when the corresponding connection counter is zero.

According to another particular feature, the memory saving operation comprises resetting the current connections that use indexed information items associated with connection counters of values lower than a threshold value. This

provision aims to free memory space although there may still be connections using the entry (indexed information item) to delete. Thanks to the resetting of these connections, their corresponding indexing tables are re-built from scratch upon the resetting, thus still enabling indexing of the deleted entries.

5           According to yet another embodiment of the invention, the method may comprise, upon receiving a request to free memory, saving the number of information items indexed in the local indexing table; flushing the global table; and indexing new information items in the local indexing table using indexes starting from the saved number. This provision makes it possible to free memory space while maintaining  
10       compatibility with the indexing mechanisms at both the encoder and decoder.

          According to another particular embodiment of the invention, the serialized binary representation of the literal values keeps said values according to their relative order in the initial list. This may be implemented when the initial order of the information items is important for the rendering by the decoder.

15           According to a fifth aspect of the invention there is provided a coding device for encoding messages, a message including an initial list of items of information, the device comprising:

          a determining means to determine a first group of indexes associated with the information items of the initial list that are already indexed in a local indexing table  
20       and a second group of literal values of the other information items of the initial list not yet indexed in said local indexing table;

          an encoder to encode the indexes of the first group;

          a compressor to binary compress at least a serialized binary representation of the literal values of the second group; and

25       a concatenation module to concatenate the first group and the second group together to obtain an encoded bitstream of the information items.

          According to a sixth aspect of the invention there is provided a decoding device for decoding encoded messages, the device comprising:

30       an obtaining means to obtain encoded data from a bitstream, the encoded data comprising a first part concatenated with a second part;

          a decoder to decode the first part to obtain a first group of indexes associated with information items already indexed in a local indexing table;

          a decompressor to binary decompress the second part to obtain a serialized binary representation of a second group of literal values corresponding to  
35       information items not yet indexed in said local indexing table; and

an output to provide a decoded list made of the information items indexed by the indexes of the first group and of the information items literally decoded from the second group.

According to a seventh aspect of the invention there is provided a coding  
5 device for encoding messages to send over a plurality of connections, a message including a list of items of information, the device comprising:

a global table shared between the connections and comprising items of information;

a plurality of local indexing tables corresponding to the plurality of  
10 connections, each local indexing table associating indexes with references to an entry of the shared global table;

an encoder to encode the list of information items of a message using the local indexing table specific to the connection to which the message belongs.

According to an eighth aspect of the invention there is provided a decoding  
15 device for decoding encoded messages received over a plurality of connections, an encoded message including an encoded list of items of information, the device comprising:

a global table shared between the connections and comprising items of information;

a plurality of local indexing tables corresponding to the plurality of  
20 connections, each local indexing table associating indexes with references to an entry of the shared global table;

a decoder to decode the encoded list of information items of an encoded  
message using the local indexing table specific to the connection to which the encoded  
25 message belongs.

Another aspect of the invention relates to a computer program product for a programmable apparatus, the computer program product comprising a sequence of instructions for carrying out one of the methods as set out above when loaded into and executed by the programmable apparatus.

Yet another aspect of the invention relates to a non-transitory computer-readable storage medium, able to be read by a programmable apparatus, storing instructions of a computer program for carrying out one of the methods as set out above when loaded into and executed by the programmable apparatus.

The coding and decoding devices, the computer program product and the  
35 computer-readable storage medium may have features and advantages that are

analogous to those set out above and below in relation to the methods of encoding or decoding, in particular that of improving the processing speed and of reducing memory consumption when encoding a list of items in a message, such as HTTP headers of HTTP messages.

5           Another aspect of the invention relates to a method of encoding substantially as herein described with reference to, and as shown in, Figure 7; Figure 9; Figures 7 and 9; Figures 7 and 11; Figures 7 and 12; Figures 7, 9 and 11; Figures 7, 9 and 12 of the accompanying drawings.

10           Another aspect of the invention relates to a method of decoding substantially as herein described with reference to, and as shown in, Figure 8; Figure 10; Figures 8 and 10; Figures 8 and 11; Figures 8 and 12; Figures 8, 10 and 11; Figures 8, 10 and 12 of the accompanying drawings.

15           At least parts of the method according to the invention may be computer implemented. Accordingly, the present invention may take the form of an entirely hardware embodiment, an entirely software embodiment (including firmware, resident software, micro-code, etc.) or an embodiment combining software and hardware aspects which may all generally be referred to herein as a "circuit", "module" or "system". Furthermore, the present invention may take the form of a computer program product embodied in any tangible medium of expression having computer usable  
20           program code embodied in the medium.

25           Since the present invention can be implemented in software, the present invention can be embodied as computer readable code for provision to a programmable apparatus on any suitable carrier medium, for example a tangible carrier medium or a transient carrier medium. A tangible carrier medium may comprise a storage medium such as a floppy disk, a CD-ROM, a hard disk drive, a magnetic tape device or a solid state memory device and the like. A transient carrier medium may include a signal such as an electrical signal, an electronic signal, an optical signal, an acoustic signal, a magnetic signal or an electromagnetic signal, e.g. a microwave or RF signal.

30           Embodiments of the invention will now be described, by way of example only, and with reference to the following drawings in which:

**Figure 1** illustrates six HTTP headers for respectively an HTTP request and an HTTP response;

35

**Figure 2** illustrates a SPDY example in which the principle of the invention is implemented;

5 **Figure 3** illustrate an example of header indexing in addition to the Deflate compression for two successive HTTP responses of a SPDY connection, according to the present invention;

10 **Figure 4** illustrates two different implementations of local indexing tables, respectively built independently and built with reference to a global table;

**Figure 5** is a plot illustrating the memory space used for storing local indexing tables according to the two approaches of **Figure 4**;

15 **Figure 6** is a block diagram illustrating components of a processing device in which embodiments of the invention may be implemented;

**Figure 7** is a flowchart illustrating general steps of an HTTP header encoding process according to one embodiment of the invention;

20 **Figure 8** is a flowchart illustrating general steps of an HTTP header decoding process according to one embodiment of the invention;

25 **Figure 9** is a flowchart illustrating general steps of an HTTP header encoding process based on a shared global table according to a second embodiment of the invention;

**Figure 10** is a flowchart illustrating general steps of an HTTP header decoding process based on a shared global table according to a second embodiment of the invention;

30 **Figure 11** is a flowchart illustrating steps of migrating current connections from a current shared global table to a new shared global table; and

**Figure 12** is a flowchart illustrating steps of optimizing a shared global table according to the age of connections.

The left part of **Figure 2** shows the conventional SPDY for a server having three connections and performing independent Deflate compression (ZIP1 , ZIP2 and ZIP3) of HTTP headers to three clients (**Figure 2a**), as well for a client having three connections to receive HTTP messages and performing independent decompression (**Figure 2c**).

The right part of **Figure 2** illustrates an embodiment of the present invention: prior to the three independent Deflate steps (ZIP1 , ZIP2 and ZIP3), header indexing is performed based on a shared global table (**Figure 2b**). Corresponding header indexing is performed at the decoder (**Figure 2d**).

This makes it possible to lighten the Deflate steps, since fewer bytes are compressed or decompressed using Deflate, thanks to the previous indexing. Furthermore, the global table shared by the connections for header indexing enables indexing memory costs to be shared also.

**Figure 3** illustrates the header indexing in addition to the Deflate compression for two successive HTTP responses of a SPDY connection. **Figure 3a** shows the situation for the first response and **Figure 3b** for the second response which occurs later on.

The four HTTP headers of each response are represented at the top of the Figure.

Regarding the first response in the SPDY connection, since no header has been already processed and indexed, no similarity is found with a previous response and thus all four HTTP headers are encoded literally with the standard SPDY encoding: the name of each HTTP header is first encoded literally, followed by its value. The processor then applies the Deflate compression to the encoded stream (i.e. a serialized binary stream) to obtain the compressed header stream that is transmitted over the network in the HTTP response.

In parallel, the four HTTP headers are stored in the list of processed headers for later indexing reuse. This list is the local indexing table having possibly (in some embodiments) references to the above-mentioned global table. In this example, the local indexing table for the SPDY connection associates each HTTP header with a specific index.

Regarding the second response which occurs later than the first response in the SPDY connection, the four HTTP headers are also processed. Three of them are exactly the same as the first three headers of the first response and are already stored in the local indexing table.

The corresponding three indexes, for example respectively 0, 1 and 2 as shown in the Figure, are retrieved from the local indexing table to encode the first three HTTP headers. This constitutes a first list or group of indexes.

5       Next, the remaining HTTP header is encoded literally using the SPDY encoding since it cannot be encoded as an index. This constitutes a second list or group of literal values, the serialized binary header stream of which is then compressed using the Deflate compression before being sent over the network in the second response, concatenated with the first list or group.

10       For the following, reference will be made to the first and second lists while a reference to groups can also be made as provided above.

Again, in parallel to sending the compressed headers, the local indexing table is updated by adding a new entry matching the literally-encoded header with a new index (index 4), so that it can be used to encode efficiently (with that index) redundancies in subsequent responses.

15       As is apparent from this example the header indexing allows processing improvements over SPDY encoding due to the reduced amount of bits that is Deflate compressed. A processing speed improvement up to twice faster can also be obtained for HTTP headers of subsequent HTTP messages, if there is high redundancy between those headers.

20       Of course, universal lossless compression techniques other than Deflate, such as LZW or BZIP2, can be used instead of the Deflate compression utilized here for illustrative purposes. Also, the invention applies to any list of information items in messages, the HTTP headers being one illustrative example.

25       In addition, due to the header indexing according to the invention which reduces the size of the binary stream to compress, the Deflate sliding window size of each connection may be reduced, thus saving both memory (less bytes to buffer) and speed (less bytes to search through). This is because the header indexing approach with a very small additional memory cost continues to improve the speed in that situation but also enables the compression ratio achieved for larger sliding window size  
30       to be maintained.

For example, when the sliding window size is made smaller than the number of bytes used to represent the HTTP headers of a message, some redundancies are not detected by the Deflate compression thus resulting in a decreased compression. To illustrate, if the same header A occurs in two successive  
35       messages, header A will be added to the current sliding window during the encoding of

the first message but may be removed from the sliding window before starting the encoding of the header A of the second message (if the window is too short). Therefore, fewer redundancies may be actually retrieved leading to a decrease in the compression efficiency.

5                This is however compensated by the header indexing since an indexed header is continuously kept in the indexing table, thus still enabling redundancies to be detected.

                 Furthermore, header indexing makes it possible to give faster access to HTTP headers at the decoder. This is because, thanks to the indexes, few bits or bytes  
10            need to be accessed from the encoded bitstream and decoded to obtain the corresponding HTTP headers (stored in the local indexing table). This is particularly true when the first list of indexes is prior to the second list of literal values in the encoded bitstream.

                 In addition to these speed and compression benefits, memory savings can  
15            be obtained by efficiently implementing the local indexing tables for several contemporaneous connections.

**Figure 4a** illustrates a first example of implementation of local indexing tables for a first connection ("#1" above) and a second connection ("#2" below). The entirety of each HTTP header encountered during the encoding of HTTP messages is  
20            added to the local indexing table specific to the connection to which the message belongs. This may result in storing the same HTTP headers twice (e.g. "Status: 200 OK") and in a total memory cost of 155 bytes.

**Figure 4b** shows a memory cost-reducing implementation of the same indexing. A global table as introduced above is used to store the HTTP headers and to  
25            associate with them references (A, B, etc.). The local indexing tables thus just associate indexes with references of the global table. This substantially reduces the memory costs, as shown by the resulting cost of 132 bytes compared to the cost of 155 bytes of **Figure 4a**.

                 The saving in memory space is about 23 bytes.

30            While references are stored within the global table in the Figure, a variant may comprise storing, within the local indexing tables, pointers to the entries of the global table, each entry only storing the HTTP headers.

**Figure 5** shows the difference in memory space used for these two approaches (independent indexing tables versus reference to a global table),  
35            depending on the number of connections. It is clear from that Figure that as the number

of connections (i.e. of local indexing tables) increases, the memory savings obtained by using the global table increase. This is because the more connections there are, the more often a single string is used by several connections and thus indexed in several local indexing tables.

5                Regardless of whether the local indexing tables is implemented independently or with references to the global table, the corresponding costs in memory for their storage is substantially compensated by the memory savings achieved by reducing the Deflate sliding window size for all the connections.

10                **Figure 6** schematically illustrates a processing device 600, either a transmitter, or a receiver, or a device embedding both functionalities, configured to implement at least one embodiment of the present invention. The processing device 600 may be a device such as a micro-computer, a workstation or a light portable device. The device 600 comprises a communication bus 613 to which there are preferably connected:

- 15                - a central processing unit 611, such as a microprocessor, denoted CPU;
- a read only memory 607, denoted ROM, for storing computer programs for implementing the invention;
- a random access memory 612, denoted RAM, for storing the executable code of the method of embodiments of the invention as well as the registers adapted to
- 20                record variables and parameters necessary for implementing methods according to embodiments of the invention; and
- a communication interface 602 connected to a communications network 603 over which digital data to be processed are transmitted.

                  Optionally, the apparatus 600 may also include the following components:

- 25                - a data storage means 604 such as a hard disk, for storing computer programs for implementing methods of one or more embodiments of the invention and data used or produced during the implementation of one or more embodiments of the invention;
- a disk drive 605 for a disk 606, the disk drive being adapted to read data from the
- 30                disk 606 or to write data onto said disk;
- a screen 609 for displaying data and/or serving as a graphical interface with the user, by means of a keyboard 610 or any other pointing means.

                  The apparatus 600 can be connected to various peripherals, such as for example a digital camera 600 or a microphone 608, each being connected to an

35                input/output card (not shown) so as to supply multimedia data to the apparatus 600.

The communication bus provides communication and interoperability between the various elements included in the apparatus 600 or connected to it. The representation of the bus is not limiting and in particular the central processing unit is operable to communicate instructions to any element of the apparatus 600 directly or  
5 by means of another element of the apparatus 600.

The disk 606 can be replaced by any information medium such as for example a compact disk (CD-ROM), rewritable or not, a ZIP disk or a memory card and, in general terms, by an information storage means that can be read by a microcomputer or by a microprocessor, integrated or not into the apparatus, possibly  
10 removable and adapted to store one or more programs whose execution enables a method according to the invention to be implemented.

The executable code may be stored either in read only memory 607, on the hard disk 604 or on a removable digital medium such as for example a disk 606 as described previously. According to a variant, the executable code of the programs can  
15 be received by means of the communication network 603, via the interface 602, in order to be stored in one of the storage means of the apparatus 600, such as the hard disk 604, before being executed.

The central processing unit 611 is adapted to control and direct the execution of the instructions or portions of software code of the program or programs  
20 according to the invention, which instructions are stored in one of the aforementioned storage means. On powering up, the program or programs that are stored in a non-volatile memory, for example on the hard disk 604 or in the read only memory 607, are transferred into the random access memory 612, which then contains the executable code of the program or programs, as well as registers for storing the variables and  
25 parameters necessary for implementing the invention.

In this embodiment, the apparatus is a programmable apparatus which uses software to implement the invention. However, alternatively, the present invention may be implemented in hardware (for example, in the form of an Application Specific Integrated Circuit or ASIC).

30 Various embodiments of encoding and decoding methods according to the invention are now described with reference to **Figures 7 to 12**. A network node may embed one of or both encoding and decoding capabilities.

If the node operates as both encoder and decoder, preference is given to the use of two distinct local indexing tables for respectively the encoding and the  
35 decoding within the same connection. This is because the headers of HTTP requests

are generally different from the headers of HTTP responses. Each local indexing table is then handled independently from the other for the same connection.

However, one skilled in the art may understand that it is still possible to use a single local indexing tables for the encoding and decoding tasks, in particular if it is  
5 observed that an large number of HTTP headers are similar within the requests and the responses of the same connection.

**Figure 7** is a flowchart illustrating general steps of an HTTP header encoding process according to one embodiment of the invention. It is recalled here that a plurality of HTTP headers have to be encoded when processing an HTTP message  
10 according to SPDY.

In the described process, only one connection is considered. Of course, the same process may be performed simultaneously by other encoders that operate for other connections. It is also made reference to the encoding of HTTP headers where the initial order of the headers is not significant and important for the decoder. When  
15 such initial order should be kept by the decoder (i.e. retrieved since the invention may possibly modify this order for the transmission), the invention may provide, within the stream, additional information (e.g. an additional header) storing the initial order of the headers. For example, this additional information may define the exact position of each indexed header within the initial order, while the other headers (i.e. without indexing)  
20 are literally encoded keeping their relative initial order.

The process starts by iterating on each HTTP header of the message to encode, as shown by the loop between steps 700 and 730.

For each header, a search is made in the local indexing table specific to the current connection to determine whether or not the header is already indexed (step  
25 710).

This search may consist in determining whether or not the HTTP header is already present in the local indexing table, when using tables as in **Figure 4a**.

When using a global table as shown in **Figure 4b**, this search comprises searching for an entry in the global table that corresponds to the header and searching,  
30 in the local indexing table, for reference to the entry found.

If the header is already indexed, its index is retrieved from the local indexing table and is stored in a first list (step 720). This first list (or group) is built to include the indexes of the message headers that are already indexed in the local indexing table.

On the contrary, when the header is not yet indexed, a first step 740 is performed to index that header in the local indexing table (possibly with reference to the global table if appropriate). Then, in a second step 750, the header is encoded literally and a compression algorithm such as Deflate is applied to the literal value (step 5 760).

Step 740 updates the local indexing table by associating, within that local indexing table, at least one literally-encoded header with a new index.

The Deflate context is kept between each processed literal value at step 760 so that the encoding result is the same as applying the Deflate compression on the 10 serialized binary representation of all the literal values determined at the occurrences of step 750. One skilled in the art would then understand that the literal values may be stored in a second list and then the Deflate compression may be performed on the serialized binary representation of that second list when exiting the loop 700-730.

Once all the headers have been processed through steps 700-750, the next 15 step consists in encoding the indexes of the first list built at the occurrences of step 720. The indexes are encoded either as single values or as ranges of values. A compression technique like Deflate can also be used. In that case, the first list and the second list can be concatenated and binary compressing may be performed on the whole serialized binary representation to compress it.

20 Once encoded, the first list of indexed headers is sent at step 770. The number of indexed headers can be typically encoded first and then each index, although a dedicated code may be defined to flag the end of the index header section in a variant.

Next the literally encoded headers are also sent at step 780.

25 For example, the first list of header indexes may be concatenated with the second list of literal values so that the resulting list is sent (after appropriate coding and compression) to perform both steps 770 and 780.

The processing savings of this process (and then of the corresponding decoding process) comes from the fact that the cost for building and searching within 30 the local indexing table is more than compensated by the savings due to not applying the Deflate compression to each indexed header.

Of course, the more headers which are indexed, the higher the processing savings. Furthermore, the added memory consumption is limited since the memory used to store the local indexing table may be compensated by decreasing the Deflate 35 sliding window size.

In the examples of **Figure 4a** and **Figure 4b** the whole HTTP header is indexed as a single entry in the local indexing table. However, other indexing strategies or decision rules may be implemented when performing the indexing step 740: for example indexing the whole header, part of the header, or not indexing the header at all depending on criteria.

According to one indexing strategy, the header names may be indexed separately from header values. In particular, header name indexing (i.e. indexing the header name) will generally be favoured over header value indexing (i.e. indexing the header value) since header names are generally more redundant than their values. For example, the 'Date' HTTP header value evolves over time and may advantageously not be indexed while the 'Date' HTTP header name may be indexed since it remains unchanged over time. Similar consideration can be used for other HTTP headers according to their nature, for example for 'Age' or 'Content-Length' headers.

When favouring the header name indexing, an indexed HTTP header of the first list may thus comprise an index indexing the HTTP name and comprise the literal HTTP value.

As one knows, memory may constitute a constraint for building the local indexing table (and global table when used). Thus, it may be decided to limit the memory size used by the indexing table, in particular by limiting the number of table entries or the memory size allocated to store the header representations. Typically, an indexing table can be limited to a buffer of 512 bytes or 1024 bytes. In that case, the decision of whether or not to index a new header (or part of it) at step 740 can be taken based on the proportion of allocated memory that is already used.

With such a limited buffer size, the additional memory cost is limited and the results are generally improved by indexing at step 740 only the header having a small binary representation, typically headers with less than 64 characters.

A refinement of that approach may comprise incrementally indexing headers with large binary representation. In other words, the header is incrementally indexed as several occurrences of that headers occur and are processed.

For example, the first time that header with large binary representation is encountered, only the first N characters (N=64 for instance) are actually stored in the indexing table and associated with an index. For that first occurrence, the remaining characters are encoded literally. The second time the same header (second occurrence) is encountered, the indexing table is updated by increasing the number of

characters to store (for instance 2N) for the same index, and so on until the whole header is stored if it occurs enough times.

According to another strategy, an indexing flag may be associated with each binary-compressed literal header to indicate that it has to be indexed in the local  
5 indexing table. The decoder can therefore be informed of any proprietary indexing strategies that it does not already know. For example, the literal encoding of the header includes a flag telling the decoder whether to add the header in the indexing table.

Proprietary strategies may cover various situations, for example the knowledge by the encoder of a precompiled list of header names to index or a  
10 precompiled list of header names not to index.

The list could for instance contain the headers that can be shared between several connections. Typically headers related to cookies may advantageously not be shared between several connections and may advantageously not be indexed. The 'Set-Cookie' header may advantageously not be indexed since its value is expected to  
15 change. In any case, headers that are not indexed will generally be efficiently compressed by the Deflate compression step 760.

Another indexing strategy may be based on the fact that some headers have a limited validity over time. This is for example the case of the "Date" header, which contains the date at which a message is sent, but with a precision to the second.

20 Since a computer is able to send many messages during one second, indexing such a date-based header can be efficient. But this will be useful only for a few messages produced within the same second. Several messages can be emitted quickly by a client to fetch all the resources relating to a Web page (style, JavaScript, images...), or by the server when responding to the client.

25 To support limited validity over time, a new header may be flagged as delta indexed, which means giving, in addition to the index, information about the difference (in second in the example) relative to the header corresponding to the index. In that case, an indexing level flag is associated with an indexed HTTP header, said indexing level flag taking, depending on the indexed HTTP header, either a first full indexing  
30 state indicating that the indexed HTTP header has the same HTTP name and HTTP value as the corresponding information item in the local indexing table, or a second partial indexing state indicating that the indexed HTTP header has the same HTTP name as the corresponding information item in the local indexing table but a different HTTP value. The difference in value may be also literally encoded and transmitted.

The "partial indexing state" thus indicates to the decoder that the header value difference has to be decoded literally.

Also as described in the above example of **Figure 7**, the indexed headers are prior to the literally encoded headers within the encoded bitstream. This provides  
5 some advantages.

For example, some headers may be notably important for network nodes to dispatch the messages to the right processing entity, like the URL value of the HTTP requests. Also, HTTP caching headers are useful for proxies. These headers are generally very redundant and are therefore quickly indexed. Providing first the indexed  
10 headers allows in most cases network nodes to get access quickly to the most important headers without having to use Deflate decompression and header string parsing techniques. Furthermore, decoding the indexed headers can be done by retrieving few bytes from the bitstream, thus making it possible to obtain a large number of decoded headers by accessing a small part of the bitstream.

However, variants can be used. For example, to limit the buffering needed, the literally-encoded headers can be sent first as a Deflate stream and then the indexed headers be sent. In that case, a specific code ends the literally-encoded header section so that it allows the decoder to quickly switch to the index-encoded header section.  
15

**Figure 7** also illustrates an example of processing where index encoding is separated from literal encoding. While this allows improved compression, it also only requires a small amount of buffering for the indexes and/or the literal values. According to a variant, the headers may be encoded and sent as soon as they are processed.  
20

Finally, a flag can be encoded as part of each literally-encoded header to indicate to the decoder that the header must be added to the index table. An alternative is that both encoder and decoder know the indexing rules exactly so that they will keep their indexing table synchronized without the need to communicate the indexing decision information.  
25

**Figure 8** is a flowchart illustrating general steps of an HTTP header decoding process according to one embodiment of the invention. In particular it is the decoding process corresponding to the encoding process of **Figure 7**.  
30

This means that the decoder operates in a similar fashion to the encoder. For example, the same indexing strategy is used by both the encoder and the decoder in such a way that they have the same local indexing tables and the same global table  
35 when used.

Similarly to the encoding process, a plurality of headers has to be processed when decoding. In the example of the Figure, the encoded indexed headers are obtained prior to the literally-encoded headers.

5       The first step consists in decoding indexed headers through steps 800 and 810.

The algorithm starts at step 800 by checking whether there is a remaining indexed header to decode. This test may be either based on decoding the number of headers or decoding a code or flag telling whether there is a remaining header to decode.

10       Each time an indexed header is to be decoded, the decoder decodes the header index from the bitstream using the local indexing table (step 810). The header is retrieved from the indexing table using its index decoded from the bitstream.

The decoded header is then added to the message header list.

15       This decoding step is fast and simple since no header parsing (separation of header name and header value, parsing of the header value in dedicated structures) is actually required. Thus header decoding speed is improved.

20       When all indexed headers have been decoded, the literally-encoded section of the bitstream is decompressed using the Deflate algorithm (step 820) and the obtained binary stream is parsed to process each literally-encoded header (steps 830 to 860).

The literal decoding starts by checking whether binary data remains to decode from the decompressed binary stream (step 830).

25       If so, the next literally-encoded header is decoded from the decompressed binary stream by retrieving the header name and its value (step 820). Once parsed, the obtained header is added to the message header list.

30       A test is then carried out to check whether this obtained header has to be indexed or not (step 850). The decoder implements the same indexing strategy as the encoder, and thus uses the same indexing rules. In the particular case where the encoder has added an indexing flag to the bitstream, that flag is searched for by the decoder in the bitstream to determine whether or not the decoded header has to be indexed.

If it has to be indexed, the decoded header is added to the local indexing table (step 860).

35       Otherwise or after this indexing step 860, following binary data to decode are considered at step 830.

Once all headers have been decoded, the full message header list is provided to an application that will process it (step 870).

As introduced above, some indexing strategies enable certain header representations to be shared between different connections of the same network node, be it decoder or encoder, through the global table. This enables a decrease in the memory cost of the local indexing tables used by the different connections at a limited processing cost: a global header table is defined for the pool of connections and may store all indexed headers. Next, for each connection, a specific indexing table stores references to the exact headers that have been indexed for that given connection.

**Figure 4b** illustrates such a strategy with a shared global table and two local indexing tables for two corresponding connections.

An encoding process based on the shared global table is now described with reference to **Figure 9**.

Contrary to the use of a shared global table, it is difficult to share the Deflate buffer between several connections since the Deflate buffers are continuously updated.

In this context, using the shared global table to share memory costs makes it possible to reduce the Deflate window size, thus achieving gains in speed and memory savings while maintaining very good compression ratio through the indexing tables.

Compared to **Figure 7**, a different approach is taken at the header indexing step 740 and at the header index retrieving steps if any (steps 710 and 720).

In detail, the encoding process starts by retrieving a header to encode (step 900).

The header is then searched for in the shared global table (step 905). Conventional search mechanisms with high search speed may be used.

If it is not found, it is encoded as a literal header (steps 920 to 940): the header is added to the shared global table (step 920) and a reference to the header representation stored in that shared global table is added to the local indexing table specific to the connection considered (the connection to which the message including the header being processed here belongs - step 930). The header is then literally encoded and Deflate compressed in the second list of literally-encoded headers (steps 935 and 940).

If the header is found in the shared global table, the associated header reference (i.e. global index in the shared global table) is retrieved in step 950. This

reference allows fast checking of whether the header has already been encoded for the connection considered, by checking the presence of that reference within the connection-specific local indexing table (step 955).

5 If the header has not already been encoded for the connection, it is indexed in the connection-specific local indexing table (step 930) and literally encoded and compressed in the second list of literally-encoded headers (steps 935 to 940).

If the header has already been encoded for the connection, the header index is retrieved from the connection-specific local indexing table and stored in the first list of indexes (step 960).

10 Once all headers have been processed, the indexes of the first list are encoded and the concatenated first and second lists are sent (steps 980 and 990) in a similar manner to steps 770 and 780. In particular, the Deflate compression of step 940 may optionally be performed only once on the serialized binary representation of all the literally-encoded headers (the second list), after step 970.

15 In a variant of this Figure, the search made at step 905 can be performed within the connection-specific local indexing table itself. In that case, the header index is directly retrieved from the connection-specific local indexing table. For example when the local indexing table contains only references to the headers which are actually stored in the shared global table using pointers to that global table, this may be done  
20 by searching, within the local indexing table, for a pointer to the searched header. In case no pointer is found, it is necessary to determine whether or not the searched header is already existing in the shared global table, so that it is inserted in that global table or not when inserted to the local indexing table. As a consequence, this variant of step 905 requires an additional search in the shared global table whenever a header is  
25 inserted in the local indexing table.

This alternative has impact on the shared global table management, in particular when there is a need to reduce the size of that global table. In that case, entries may not be easily removed from the global table. One possibility is to switch from one global table to another and remove the old global table at some point as  
30 explained below with reference to the global table migration of **Figure 11** or to keep the shared global table and the local indexing tables tightly synchronized as explained below with reference to the global table optimization of **Figure 12**.

In another variant, the shared global table may be managed so that only shareable headers are actually stored in the global table. This is because memory

savings are obtained only if the headers in the global table are reused over several connections.

In particular, only frequent shareable headers can be actually stored in that global table. The shareable header may for example be predetermined and known by  
5 both the decoder and the encoder based on the HTTP standard. They may also be headers other than the below defined private headers. This makes it possible to optimize memory use with the headers that are worthy to be stored in a shareable manner.

Various level of management may be applied for the other headers (i.e. not determined as being shared).

10 For example, private headers (such as cookies or sensitive data, connection-specific data such as "Host" or "Referer", etc.) may be literally-encoded and then Deflate compressed.

In another example, a two-level indexing mechanism is used, whereby frequent and shared headers are stored in the shared global table while frequent but  
15 not shared (i.e. private) headers are stored in the connection-specific local indexing table. The other headers are not indexed. In that case, the local indexing table comprises headers indexed by reference to the global table and headers indexed within the local indexing table alone. Furthermore, deciding to index a header in the local indexing table alone or by reference to the global table depends on a determined  
20 frequency of that header.

The private, shareable and/or frequent headers may be defined as such in predefined lists provided to the encoder (and if needed to the decoder).

**Figure 10** is a flowchart illustrating the corresponding decoding based on the shared global table. As for **Figure 9** which is an adaptation of **Figure 7**, the process  
25 of **Figure 10** is also an adaptation of **Figure 8**.

However the process of **Figure 10** requires more processing than the one of **Figure 8** since for every newly indexed header, an additional search is performed in the global table, which can be expensive at the decoder. This is in contrast to the processing at the encoder where only one header search is required with or without the  
30 use of the shared global table. This allows memory saving at the expense of processing speed. To mitigate these drawbacks, the indexing criteria may be limited to only index some headers, for example the shareable headers and/or those that are the most frequent in several connections as provided above.

The indexed headers are first decoded since they are generally put first in  
35 the bitstream (steps 1000 and 1010). The decoder decodes the header index from the

bitstream. It then either finds the header directly in the connection-specific local indexing table thanks to the decoded header index, or finds the header using the global table and the reference associated with the decoded header index in the connection-specific local indexing table. The decoded header is then stored in the message header list.

5                This is because the connection-specific local indexing table can either store the header representation or a reference to its representation that is stored in the global table. The global table may directly store headers as strings or in parsed representations.

Literally-encoded headers are then decoded through steps 1020 to 1090.

10              The bitstream is first decompressed using Deflate (step 1020) to obtain a binary stream.

The decompressed binary stream is then processed to parse it until no more data is available (test 1030). As long as binary data is available, the header is first parsed (the header name and header value are obtained) and added to the message header list (step 1050).

15

Similarly to step 850, the strategy for indexing the header is determined at step 1060.

If the header has not to be indexed, the decoding process continues with another header (loop back to step 1030).

20              Otherwise, it is tested whether the header is already in the global table at step 1070.

In the negative, the header is added to the global table as a new entry, at step 1080, meaning that a reference associated with that new header is created.

25              In the affirmative or after step 1080, the header is added to the connection-specific local indexing table at step 1090, by association of a corresponding index with the reference to the global table header entry.

When all the binary data have been processed, the decoded headers (i.e. the whole message header list) are provided as an output to, for instance, an application (step 1040).

30              In particular embodiments of the invention, described below with reference to **Figures 11 and 12**, optimizations of the shared global table use are provided with a goal of reducing memory consumption.

Since the shared global table will span the lifetime of several connections, some headers therein may become irrelevant when a given connection is closed.

The shared global table will continuously grow while the number of headers used may not actually grow. There is therefore a need to ensure that headers in the shared global table remain relevant to existing connections and to remove from the global table headers that are no longer relevant to current connections. The  
5 optimizations of **Figures 11** and **12** address this issue at the encoder. Of course, similar optimization should be provided at the corresponding decoder.

**Figure 11** is a flowchart illustrating steps of migrating current connections from a current global table to a new global table. The new global table will be devoid of non-used headers unlike the current global table.

10 This process is based on switching from one shared global table to another at a given time: at some point a shared global table is created and all new connections will use the new shared global table while the old connections will remain with the initial shared global table. The old shared global table will have the number of its connections decreasing to a small number, at which time the remaining connections will be  
15 migrated to the new shared global table. The old shared global table will be deleted to free memory space.

This approach ensures that the overall memory used by the global table does not continuously grow over time.

The process of **Figure 11** starts at step 1100 by waiting for a network event  
20 received by the network node encoding messages based on local indexing tables and a current shared global table. It is assumed that all the current connections are linked to the same current shared global table A, meaning that their connection-specific local indexing tables all refer to that global table A.

If the incoming network event is a new connection (test at step 1110),  
25 whether the current shared global table needs to be refreshed or not is first checked (step 1115). Refresh criteria may be based on different information: current global table age, and/or number of connections closed since the global table has been created, etc. Furthermore, if the current global table is already in a migration phase (meaning that a new global table has already been created, as explained below), there is no need to  
30 refresh it.

If the global table needs to be refreshed, a new global table B is created at step 1120. That means that the new global table is created upon receiving a request for a new connection, depending on at least one criterion relating to the currently shared global table. This global table B can be initialized as an empty table or can be prefilled  
35 with commonly used headers. The new connection (corresponding to the event

received) is then linked to that new global table at step 1125, meaning that the corresponding connection-specific local indexing table will associate indexes with references to that new global table B (and not to the current global table A).

5 When the received network event is either a new connection while a migration is in course (test at step 1110) or an end of connection (test at step 1130), whether a header needs to be processed or not is checked at step 1150. A migration is in course when two global tables A and B exist at the same time.

10 If no header needs to be processed (output "no" at step 1150), a processing dedicated to the received network event is performed at step 1175 before waiting for a new network event by looping back to step 1100.

15 If a header has to be processed, that header is encoded based on the connection-specific local indexing table (an index is retrieved if it already exists, using the global table to which the connection is linked) and the linked global table and connection-specific local indexing table are updated if the processed header needs to be inserted (step 1155).

20 An optional step (step 1160) may be carried out to cap the total memory used by the global tables A and B: if the total memory is above a given threshold once the processed header has been indexed and added to the tables (test 1160), some entries may be removed from the global table A (step 1170) to decrease memory consumption. Examples of entries that can be removed are the oldest entries, the least frequently used entries, etc. Of course, when an entry is removed from the global table, the corresponding entry in the local indexing tables also has to be removed since the reference to the global table no longer exists.

25 The process then loops back to step 1100 of waiting for a next network event.

When the received network event is an end of connection event and the migration is in course (detected at step 1130), whether only a few connections are linked to the current global table A or not is checked at step 1135. Few may advantageously mean a predetermined number of connections, for example less than 5% of on-going connections.

30 If this is the case, all those connections (linked to global table A) are migrated to the new global table B (step 1140): this comprises removing entries used by the current connections from the current global table A; adding those removed entries into the new global table B; and then deleting said currently shared global table from memory. Of course the local indexing tables of those migrated connections also  
35 have to be updated so that they now refer to the migrated entries in the new global

table B, instead of the deleted entries of the deleted global table A. Adding the removed entries to B may be performed while avoiding duplication of those entries. That is, if a removed entry already exists in B, it is not added thereto but the local indexing tables of the migrated connections are updated so that they now refer to this  
5 already existing entry.

In this part of the process, migrating current connections is thus triggered upon receiving a request to end a current connection.

The advantage of this process is that the migration phase processing is limited in time and may be carried out when the network node has processing  
10 capacities available. In particular, step 1140 may be carried out asynchronously.

A different approach is now described with reference to **Figure 12**, where instead of creating and deleting global tables, the current global table is kept synchronized with the current connections: all the headers in the global table are used by the current connections (i.e. referenced by the local indexing tables). Such an  
15 approach is preferable to the approach of **Figure 11**, when there is a risk of the current global table not remaining naturally synchronized with the connections over a long period of time.

In the approach of **Figure 12**, each indexed header in the global table is associated with a connection counter that counts the number of connections using that  
20 indexed header, and a memory saving operation is carried out based on connection counters of the indexed headers. In particular, the memory saving operation comprises removing any indexed header from the global table whose corresponding connection counter is zero. This ensures the global table only stores headers that are used by the current connections. The global table is thus synchronized with the current  
25 connections' state.

The process may start at step 1235 by waiting for a network event received by the network node.

Based on the received event, whether a new header has to be processed or not is first checked at step 1200.

30 In the affirmative, whether the new header should be indexed or not is checked at step 1210.

If so, whether its representation is already in the global table or not is checked at step 1220.

If it is already present, a counter associated with the entry of that header in  
35 the global table is incremented at step 1230. The counter is used to reflect the number

of connections that use the corresponding header. In this context, it is named "connection counter".

If it is not present, the header is added to the global table and its connection counter is initialized to 0 (step 1225) before being incremented by one (step 1230).

5           After step 1230, the process loops back to step 1235 of waiting a next network event.

If the received event is not a new header to process, it is checked whether or not it is an end-of-connection event (test at step 1240).

10           If it is, a connection has thus stopped (meaning that the corresponding local indexing table will be deleted in the near future). All the headers indexed in the local indexing table specific to the stopped connection are processed before deletion.

For each entry in that local indexing table (test 1250), the corresponding connection counter in the global table is decremented by one at step 1255. The "corresponding connection counter" is the counter associated, in the global table, with the  
15           entry having the reference indicated in the entry considered of the local indexing table.

If the connection counter reaches 0 (test 1260), the header representation is removed from the global table at step 1265 to free memory space.

Once all the entries of the local indexing table have been processed, the process loops back to step 1235.

20           If the received network event is neither a header to process nor an end-of-connection event, it is conventionally processed by dedicated processing at step 1270 before looping back to step 1235.

The process of **Figure 12** allows tight synchronization between the global table and the local indexing tables (i.e. the current connections). This comes at the  
25           expense of the processing that must take place synchronously during the network communications and some additional memory consumption, since a connection counter is kept up to date for each header of the global table.

In addition to these processes of optimizing the global table use, there may be additional optimization processes.

30           For example, it may be necessary for the network node to free memory space. In that case, the global table may be scanned to find appropriate entries to delete and thus to reduce the global table size. One approach comprises resetting the current connections that use indexed headers associated with connection counters of values lower than a threshold value. The same connection counters as defined above  
35           are used. Upon determining all the headers for which the connection counters have a

low value, the corresponding connections (those using the determined headers) are reset. This allows removal of all the low-counter headers from the global table at the expense of forcing the client to open a new fresh connection that has no indexed headers.

5                   Of course criteria other than the connection counter alone may be used to forcefully remove headers from the global table. For example, a criterion may be based on a creation date or a last usage date (the date doesn't need to be a full date, but can be a short integer increased regularly, as for example a 4-bit integer corresponding to 1 minute spans, enabling times from 0 to 15 minutes to be measured) defined for each  
10 entry in the global table. Thus, only the oldest headers are removed.

                  Another additional optimization process may comprise, upon receiving a request to free memory, saving, in memory, the number of headers indexed in the local indexing table (i.e. generally the maximum number of indexes already used) ; deleting the global table; and indexing new headers in the local indexing table using indexes  
15 starting from the saved number. This may be of course carried for all the local indexing tables corresponding to connections linked to the deleted global table.

                  Thus each local indexing table keeps the number of past entries and will index new entries (i.e. headers) using that number as the first index. This allows new header indexes to remain compatible with the coder indexing tables.

20                   In particular, the decoder is not forced to perform the same deletion of global table since it can detect new indexing starting from the number of past entries. However, the decoder can implement the same deletion strategy as the coder, meaning that they both receive the request to free memory at the same time and perform the same deletion.

25                   In a variant, the decoder may add a flag to the bitstream to indicate to the decoder when it has performed memory freeing. This enables the decoder to also flush its memory (global table) at the same point in the decoding process.

                  The above processes are based on indexing information items (HTTP headers in the case of SPDY).

30                   A processor may however decide to enable or disable this header indexing for any or all network connections. Criteria on which to base this decision may relate to speed or compression improvements.

                  As regards compression improvements, the header indexing approach (i.e. use of indexes) may be enabled based on the computation of Deflate compression ratios.

For example, the compression ratio achieved for the headers of the first and second requests of a current connection may be computed. Next, if the compression ratio for the second request (ratio H2) is not much higher than the compression ratio of the first request (ratio H1), the header indexing provides substantial compression improvements. It has been observed that 80% is a good threshold for the ratio H2/H1 to trigger header indexing (if  $H2/H1 \geq 80\%$ ).

This is may be because the Deflate compression context built from the first request appears not to be efficiently used for the second request, probably because the context (sliding window) is too small.

As regards speed improvements, the number of messages exchanged in a connection appears to be closely related to speed performance.

To illustrate this, when a single message is exchanged in a connection, the cost of building a local indexing table, even if small, is not compensated by the use of that indexing table.

As soon as an increasing number of messages is exchanged, the indexing table is increasingly used, leading to speed gains. Typically, header indexing provides substantial speed improvements for exchanges with more than five messages and 2x speed improvements for exchanges with more than ten messages.

A network node may then decide to compute statistics on the number of messages per connection, and based on an average number for all networked nodes or an average number per network node, it may be decided to use the header indexing or not.

Although the present invention has been described hereinabove with reference to specific embodiments, the present invention is not limited to the specific embodiments, and modifications which lie within the scope of the present invention will be apparent to a person skilled in the art. Many further modifications and variations will suggest themselves to those versed in the art upon making reference to the foregoing illustrative embodiments, which are given by way of example only and which are not intended to limit the scope of the invention as determined by the appended claims. In particular different features from different embodiments may be interchanged, where appropriate.

CLAIMS

1. A method of encoding messages, a message including an initial list of items of information, the method comprising:

5                   determining a first group of indexes associated with the information items of the initial list that are already indexed in a local indexing table and a second group of literal values of the other information items of the initial list not yet indexed in said local indexing table;

                  encoding the indexes of the first group;

10                  binary compressing at least a serialized binary representation of the literal values of the second group; and

                  concatenating the first group and the second group together to obtain an encoded bitstream of the information items.

15                  2. The method of Claim 1, wherein the first group of indexes is prior to the second group of literal values in the encoded bitstream.

                  3. The method of Claim 1, wherein the step of concatenating generates a serialized binary representation of the concatenated first and second groups, and the step of binary compressing compresses the whole serialized binary representation.

20                  4. The method of Claim 1, wherein the items of information are HTTP headers of an HTTP message.

                  5. The method of the Claim 4, wherein an HTTP header is made of at least an HTTP name and an HTTP value, and an indexed HTTP header of the first group comprises an index indexing the HTTP name and comprises the literal HTTP value.

25                  6. The method of Claim 4, wherein an HTTP header is made of at least an HTTP name and an HTTP value, and the method further comprises associating an indexing level flag with an indexed HTTP header of the first group, said indexing level flag taking, depending on the indexed HTTP header, either a first full indexing state indicating that the indexed HTTP header has the same HTTP name and HTTP value as the corresponding information item in the local indexing table, or a second partial indexing state indicating that the indexed HTTP header has the same HTTP name as  
30                  the corresponding information item in the local indexing table but a different HTTP value.

35                  7. The method of Claim 1, further comprising updating the local indexing table by associating, within that local indexing table, at least one information item of the second group with a new index.

8. The method of Claim 7, further comprising associating, in the encoded stream, an indexing flag with each binary-compressed information item of the second group that is associated with a new index during said updating.

5 9. The method of Claim 7, further comprising determining a frequency of occurrence of an item of information of the second group, and deciding to index said item of information in the local indexing table depending on the determined frequency.

10. The method of Claim 1, wherein the local indexing table associates each index with a reference to an entry of a global table, said entry of the global table comprising the item of information associated with said index.

10 11. The method of Claim 10, wherein the local indexing table comprises at least one item of information indexed by reference to the global table and at least one item of information indexed within the local indexing table alone.

12. The method of Claim 11, further comprising determining a frequency of occurrence of an item of information, and deciding to index said item of information in  
15 the local indexing table alone or by reference to the global table depending on the determined frequency.

13. The method of Claim 10, wherein the global table is shared between a plurality of compressors and comprises items of information from messages processed by said plurality of compressors, each compressor having its local indexing table to  
20 index, by reference to the shared global table, items of information from messages that it processes.

14. The method of Claim 10, wherein determining the first and second groups comprises searching for entries in the global table that correspond to the information items of the initial list and searching, in the local indexing table, for  
25 references to the entries found.

15. The method of Claim 10, wherein the global table is shared between a plurality of network connections and referred to by a corresponding plurality of respective local indexing tables, and

the method further comprising, at some point during the encoding,  
30 migrating current connections from said shared global table to a new global table.

16. The method of Claim 15, wherein the new global table is created upon receiving a request for a new connection, depending on at least one criterion relating to the currently shared global table.

17. The method of Claim 15, further comprising removing entries from the currently shared global table based on the amount of memory used to store both currently shared and new global tables.

18. The method of Claim 15, wherein migrating current connections is  
5 triggered upon receiving a request to end a current connection.

19. The method of Claim 15, wherein migrating current connections comprises removing entries used by the current connections from the currently shared global table; adding those removed entries to the new global table; and then deleting said currently shared global table from memory.

10 20. The method of Claim 10, wherein the global table is shared between a plurality of network connections and referred to by a corresponding plurality of respective local indexing tables, each indexed information item in the shared global table being associated with a connection counter that counts the number of connections using that indexed information item, and

15 the method further comprising performing a memory saving operation based on connection counters of the indexed information items.

21. The method of Claim 20, wherein the memory saving operation comprises removing the indexed information item from the shared global table when the corresponding connection counter is zero.

20 22. The method of Claim 20, wherein the memory saving operation comprises resetting the current connections that use indexed information items associated with connection counters of values lower than a threshold value.

23. The method of Claim 10, further comprising upon receiving a request to free memory, saving the number of information items indexed in the local indexing  
25 table; flushing the global table; and indexing new information items in the local indexing table using indexes starting from the saved number.

24. The method of Claim 1, further comprising choosing between a first encoding of the initial list of information items using said steps of determining, concatenating and binary compressing, and a second encoding of the initial list of  
30 information items using only binary compression of a serialized binary representation of the item literal values;

a sliding binary window used for the binary compression of the first encoding being of lower size than a sliding binary window used for the binary compression of the second encoding.

25. The method of Claim 1, wherein the serialized binary representation of the literal values keeps said values according to their relative order in the initial list.

26. A method of decoding encoded messages, the method comprising:  
obtaining encoded data from a bitstream, the encoded data comprising a  
5 first part concatenated with a second part;  
decoding the first part to obtain a first group of indexes associated with information items already indexed in a local indexing table;  
binary decompressing the second part to obtain a serialized binary representation of a second group of literal values corresponding to information items  
10 not yet indexed in said local indexing table; and  
providing a decoded list made of the information items indexed by the indexes of the first group and of the information items literally decoded from the second group.

27. A method of encoding messages to send over a plurality of connections,  
15 a message including a list of items of information, the method comprising:  
encoding the list of information items of a message using a local indexing table specific to the connection to which the message belongs;  
wherein a global table is shared between the connections and comprises items of information; and the local indexing table of each connection associates  
20 indexes with references to an entry of the shared global table.

28. A method of decoding encoded messages received over a plurality of connections, an encoded message including an encoded list of items of information, the method comprising:

25 decoding the encoded list of information items of an encoded message using a local indexing table specific to the connection to which the encoded message belongs;

wherein a global table is shared by all the connections and comprises items of information; and the local indexing table of each connection associates indexes with references to an entry of the shared global table.

30 29. A coding device for encoding messages, a message including an initial list of items of information, the device comprising:

a determining means to determine a first group of grouping indexes associated with the information items of the initial list that are already indexed in a local indexing table and a second group of literal values of the other information items of the  
35 initial list not yet indexed in said local indexing table;

an encoder to encode the indexes of the first group;  
a compressor to binary compress at least a serialized binary representation of the literal values of the second group; and

5 a concatenation module to concatenate the first group and the second group together to obtain an encoded bitstream of the information items.

**30.** A decoding device for decoding encoded messages, the device comprising:

an obtaining means to obtain encoded data from a bitstream, the encoded data comprising a first part concatenated with a second part;

10 a decoder to decode the first part to obtain a first group of indexes associated with information items already indexed in a local indexing table;

a decompressor to binary decompress the second part to obtain a serialized binary representation of a second group of literal values corresponding to information items not yet indexed in said local indexing table; and

15 an output to provide a decoded list made of the information items indexed by the indexes of the first group and of the information items literally decoded from the second group.

**31.** A coding device for encoding messages to send over a plurality of connections, a message including a list of items of information, the device comprising:

20 a global table shared between the connections and comprising items of information;

a plurality of local indexing tables corresponding to the plurality of connections, each local indexing table associating indexes with references to an entry of the shared global table;

25 an encoder to encode the list of information items of a message using the local indexing table specific to the connection to which the message belongs.

**32.** A decoding device for decoding encoded messages received over a plurality of connections, an encoded message including an encoded list of items of information, the device comprising:

30 a global table shared between the connections and comprising items of information;

a plurality of local indexing tables corresponding to the plurality of connections, each local indexing table associating indexes with references to an entry of the shared global table;

a decoder to decode the encoded list of information items of an encoded message using the local indexing table specific to the connection to which the encoded message belongs.

- 5        33. A computer program product for a programmable apparatus, the computer program product comprising a sequence of instructions for carrying out the method according to any one of Claims 1 to 28 when loaded into and executed by the programmable apparatus.

- 10       34. A non-transitory computer-readable storage medium storing instructions of a computer program for carrying out the method according to any one of Claims 1 to 28.

35. A method of encoding substantially as herein described with reference to, and as shown in, Figure 7; Figure 9; Figures 7 and 9; Figures 7 and 11; Figures 7 and 12; Figures 7, 9 and 11; Figures 7, 9 and 12 of the accompanying drawings.

- 15       36. A method of decoding substantially as herein described with reference to, and as shown in, Figure 8; Figure 10; Figures 8 and 10; Figures 8 and 11; Figures 8 and 12; Figures 8, 10 and 11; Figures 8, 10 and 12 of the accompanying drawings.

1/9

## HTTP headers in HTTP request

**Host:** en.wikipedia.org  
**User-Agent:** Mozilla/5.0 (Windows; en-US) Gecko/20110303 Firefox/3.6.15  
**Accept:** text/html,application/xhtml+xml,application/xml;q=0.9,\*/\*;q=0.8  
**Accept-Language:** en-us,en;q=0.5  
**Keep-Alive:** 115  
**Proxy-Connection:** keep-alive

## HTTP headers in HTTP response

**Date:** Thu, 15 Sep 2011 13:38:15 GMT  
**Server:** Apache  
**Location:** http://en.wikipedia.org/wiki/SPDY  
**Content-Encoding:** gzip  
**Content-Type:** text/html; charset=utf-8  
**Content-Length:** 20

Fig. 1




Fig. 2a




Fig. 2b




Fig. 2c




Fig. 2d

2/9

(3a)

First HTTP  
response

|   |                                        |
|---|----------------------------------------|
| 0 | Status: 200                            |
| 1 | Version: HTTP/1.1                      |
| 2 | Server: Apache                         |
| 3 | Date: Fri, 21 Jan 2011 08:54:00<br>GMT |



|                                                                                                                                                                                                                            |  |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--|
| <b>Similarities</b> ( <i>list of indexes</i> )<br>None ( <i>since no headers yet encoded</i> )                                                                                                                             |  |
| <b>Differences</b> ( <i>SPDY encoding</i> )<br>Number of headers: 4<br>Name length: 6 – Name: "status"<br>Value length: 3 – Value: "200"<br>Name length: 7 – Name: "version"<br>Value length: 8 – Value: "HTTP/1.1"<br>... |  |

Update  
listzlib  
deflateCompressed  
Headers #1

| List of Processed Headers |                                        |
|---------------------------|----------------------------------------|
| <i>Index</i>              | <i>HTTP header</i>                     |
| 0                         | Status: 200                            |
| 1                         | Version: HTTP/1.1                      |
| 2                         | Server: Apache                         |
| 3                         | Date: Fri, 21 Jan 2011 08:54:00<br>GMT |

(3b)

Second HTTP  
response

|   |                                        |
|---|----------------------------------------|
| 0 | Status: 200                            |
| 1 | Version: HTTP/1.1                      |
| 2 | Server: Apache                         |
| 3 | Date: Sat, 19 Feb 2011<br>12:23:00 GMT |



|                                                                                                                                                                        |  |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--|
| <b>Similarities</b> ( <i>list of indexes</i> )<br>Headers 0 to 2                                                                                                       |  |
| <b>Differences</b> ( <i>SPDY encoding</i> )<br>Number of headers: 1<br>Name length: 4 – Name: "Date"<br>Value length: 29 – Value: "Sat, 19 Feb<br>2011 12:23:00 GMT »" |  |

Update  
listzlib  
deflateCompressed  
Headers #2

| List of Processed Headers |                                        |
|---------------------------|----------------------------------------|
| <i>Index</i>              | <i>HTTP header</i>                     |
| 0                         | Status: 200                            |
| 1                         | Version: HTTP/1.1                      |
| 2                         | Server: Apache                         |
| 3                         | Date: Fri, 21 Jan 2011 08:54:00<br>GMT |
| 4                         | Date: Sat, 19 Feb 2011 12:23:00<br>GMT |

Fig. 3

3/9

| Connection #1 |                                |             |
|---------------|--------------------------------|-------------|
| <i>Index</i>  | <i>HTTP header</i>             | <i>Cost</i> |
| 0             | Status: 200 OK                 | 12          |
| 1             | Version: HTTP/1.1              | 15          |
| 2             | Server: Apache/2.2.14 (Ubuntu) | 28          |
| 3             | Content-Type: text/html        | 22          |

| Connection #2 |                                |             |
|---------------|--------------------------------|-------------|
| <i>Index</i>  | <i>HTTP header</i>             | <i>Cost</i> |
| 0             | Status: 200 OK                 | 12          |
| 1             | Version: HTTP/1.1              | 15          |
| 2             | Server: Apache/2.2.14 (Ubuntu) | 28          |
| 3             | Content-Type: image/jpeg       | 23          |

|              |  |            |
|--------------|--|------------|
| <b>TOTAL</b> |  | <b>155</b> |
|--------------|--|------------|

Fig. 4a

| Connection #1 |                  |             | Connection #2 |                  |             |
|---------------|------------------|-------------|---------------|------------------|-------------|
| <i>Index</i>  | <i>Reference</i> | <i>Cost</i> | <i>Index</i>  | <i>Reference</i> | <i>Cost</i> |
| 0             | A                | 4           | 0             | A                | 4           |
| 1             | B                | 4           | 1             | B                | 4           |
| 2             | C                | 4           | 2             | C                | 4           |
| 3             | D                | 4           | 3             | E                | 4           |

| Server Global List of Processed Headers |                                |             |
|-----------------------------------------|--------------------------------|-------------|
| <i>Ref.</i>                             | <i>HTTP header</i>             | <i>Cost</i> |
| A                                       | Status: 200 OK                 | 12          |
| B                                       | Version: HTTP/1.1              | 15          |
| C                                       | Server: Apache/2.2.14 (Ubuntu) | 28          |
| D                                       | Content-Type: text/html        | 22          |
| E                                       | Content-Type: image/jpeg       | 23          |
| <b>TOTAL</b>                            |                                | <b>132</b>  |

Fig. 4b

4/9




Fig. 5






Fig. 7

5/9



6/9




Fig. 9

7/9




Fig. 10




Fig. 11




Fig. 12

# INTERNATIONAL SEARCH REPORT

International application No  
PCT/IB2011/055441

A. CLASSIFICATION OF SUBJECT MATTER  
INV. H03M7/30 H04L29/08  
ADD.

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

## B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)  
H03M H04L

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. |
|-----------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------|
| X         | <p>Jon Butler ET AL: "A Proposal for Shared Dictionary Compression over HTTP" ,<br/>'<br/>8 September 2008 (2008-09-08) , pages 1-18,<br/>XP55026339,<br/>Retrieved from the Internet:<br/>URL: http://www.bl ogs .zeenor.com/wp-content<br/>/uploads/2011/01/Shared_Di cti onary_Compres<br/>si on_over_HTTP.pdf<br/>[retrie ved on 2012-05-07]<br/>page 1 - page 2</p> <p style="text-align: center;">-----<br/>- / - -</p> | 1-36                  |



Further documents are listed in the continuation of Box C.



See patent family annex.

### \* Special categories of cited documents :

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

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

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

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

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

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

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

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

"&" document member of the same patent family

Date of the actual completion of the international search

9 May 2012

Date of mailing of the international search report

16/05/2012

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

Nowbakht Irani , Al i

# INTERNATIONAL SEARCH REPORT

International application No

PCT/IB2011/055441

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

| Category* | Citation of document, with indication, where appropriate, of the relevant passages                                                                                                                                                                                                                                                                                                                                        | Relevant to claim No. |
|-----------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------|
| A         | <p>MUN CHOON CHAN ET AL: "Cache-based compaction: a new technique for optimizing Web transfer" ,<br/> INFOCOM '99. EIGHTEENTH ANNUAL JOINT CONFERENCE OF THE IEEE COMPUTER AND COMMUNICATIONS SOCIETIES. PROCEEDINGS. IEEE NEW YORK, NY, USA 21-25 MARCH 1999, PISCATAWAY, NJ, USA, IEEE, US, vol . 1, 21 March 1999 (1999-03-21) , pages 117-125, XP010323762 ,<br/> ISBN: 978-0-7803-5417-3<br/> the whole document</p> | 1-36                  |
| A         | <p>wo 2011/067769 AI (INFOGIN LTD [IL] ; YALOZ ARI EL [IL] ; SHTERENZON ROMAN [IL] )<br/> 9 June 2011 (2011-06-09)<br/> the whole document</p>                                                                                                                                                                                                                                                                            | 1-36                  |

## INTERNATIONAL SEARCH REPORT

### Information on patent family members

International application No

PCT/IB2011/055441

| Patent document<br>cited in search report | Publication<br>date | Patent family<br>member(s) | Publication<br>date |
|-------------------------------------------|---------------------|----------------------------|---------------------|
| W0 2011067769                             | A 1                 | 09-06-2011                 | NONE                |
| -----                                     |                     |                            |                     |