

UK Patent Application (19) GB (11) 2 339 368 (13) A

(43) Date of A Publication 19.01.2000

(21) Application No 9814669.9

(22) Date of Filing 08.07.1998

(71) Applicant(s)

International Business Machines Corporation
(Incorporated in USA - New York)
Armonk, New York 10504, United States of America

(72) Inventor(s)

Paul A Beaven
Davin Morris Coleman
Brian Phillip Ernest Fisk
Barry Keith Aldred

(74) Agent and/or Address for Service

Kevin J Fournier
IBM United Kingdom Limited, Intellectual Property
Department, Hursley Park, WINCHESTER, Hampshire,
SO21 2JN, United Kingdom

(51) INT CL⁷
H04L 12/56

(52) UK CL (Edition R)
H4P PPS

(56) Documents Cited
GB 2314740 A

(58) Field of Search
UK CL (Edition Q) H4P PFD PPEC PPS
INT CL⁶ H04L 12/08 12/56
Online : EPODOC, WPI

(54) Abstract Title

Data communications protocol with efficient packing of datagrams

(57) A data communications apparatus connected to a network (e.g. the Internet) through which datagrams flow from the data communications apparatus to at least one other data communications apparatus also connected to the network, the apparatus having: a means for receiving data which is to be transmitted over the network, such data being broken up into data records DR1-DR8 and a means for packing the data records into datagrams 1-3 for transmission over the network, including a means for assigning data records to datagrams; wherein the means for packing takes the data records out of sequential order when assigning data records to datagrams.




FIG. 3

GB 2 339 368 A





FIG. 1

FIG. 2




FIG. 3

DATA COMMUNICATIONS APPARATUS, METHOD AND COMPUTER
PROGRAM PRODUCT IMPLEMENTING DATA COMMUNICATIONS
PROTOCOL WITH EFFICIENT PACKING OF DATAGRAMS

5

Field of the Invention

10 The present invention relates to protocols (i.e., sets of rules) for the communication of digital data between a sender and receiver (e.g., the sender and receiver are computer systems located at different physical locations).

15 Background of the Invention

15 The usual form of digital data communication involves sending such digital data in the form of datagrams over a network. See, e.g.,
20 *Computer Networks* by Andrew S. Tanenbaum, Prentice-Hall Inc., Englewood Cliffs, NJ 1981. A datagram is a group of bits representing a portion of the digital data to be transmitted and each datagram is treated individually by the network. Thus, if a large amount of data is to be transmitted from a sender data processing unit to a receiver data processing unit, the sender data processing unit breaks up the large amount of data into smaller sized blocks (called data records
25 hereinbelow) and packs these records into datagrams for transmission over a data communications network to the receiver data processing unit. The datagrams which are being sent over the network range in size from 1 to some fixed number (say n) bytes. The data records which are to be packed into the datagrams also may range in size from 1 to some other fixed
30 number (say m) bytes.

35 Whenever a datagram is sent over the network, there are various overheads associated with such transmission. Firstly, it takes time to process each datagram at the sending and receiving ends. This time is split into two components: a fixed amount per each datagram and an amount proportional to the size of the datagram. The fixed amount is usually significantly greater than the variable amount. Secondly, each datagram must have extra data associated with it (address information etc.) which involves a fixed amount of extra data to be added to each datagram. Due
40 to such overheads, it is therefore inefficient to send small datagrams over the network.

5 The usual technique for packing data records into datagrams involves taking the data records in sequential order. For example, a list of variable length data records of lengths 200 bytes, 250 bytes, 100 bytes and 50 bytes are to be packed into a plurality of 512-byte datagrams. The first two records can be packed into the first datagram (because 200 bytes plus 250 bytes will fit in a 512 byte datagram). However, the next data record (of 100 bytes) will not fit (because 450 plus 100 will be greater than the 512 byte capacity of a datagram). Thus, the 100 byte record will have to be packed into the next datagram.

10

15 The problem with this prior technique, however, is that there is some space left over in the first datagram which has not been utilized. Thus, the overheads discussed above are incurred without obtaining as much data transmission as would otherwise be possible. Further, when certain communications networks are used, such as mobile networks, transmission time can be very expensive, so it is highly advantageous to flow as few datagrams as possible over the network with as much data as possible packed into each datagram. Thus, the present state of the art in this area results in an inefficient and potentially very expensive systems architecture.

20

Summary of the Invention

25 According to a first aspect, the invention provides a data communications apparatus connected to a network through which datagrams flow from the data communications apparatus to at least one other data communications apparatus also connected to the network, the apparatus having: a means for receiving data which is to be transmitted over the network, such data being broken up into data records; and a means for packing the data records into datagrams for transmission over the network, including a means for assigning data records to datagrams; wherein the means for packing takes the data records out of sequential order when assigning data records to datagrams.

35

30 Preferably, the means for packing assigns, to a datagram presently being packed, the largest unassigned data record that will fit in the datagram.

40

35 Preferably, the means for assigning considers only one datagram at a time when assigning data records to datagrams. However, alternatively,

Preferably, the means for assigning considers only one datagram at a time when assigning data records to datagrams. However, alternatively, the means for assigning considers a plurality of datagrams at a time when assigning data records to datagrams.

5

Preferably, the size of each data record must be less than or equal to the size of each datagram. However, alternatively, the size of a data record is larger than the size of a datagram.

10

According to a second aspect, the invention provides a method of carrying out the functions recited with respect to the first aspect.

15

According to a third aspect, the invention provides a computer program product stored on a computer readable storage medium, for when running on a computing system, carrying out the functions recited with respect to the first aspect of the invention.

20

Thus, by taking data records out of sequential order when assigning data records to datagrams and thereby packing as much data as possible into each datagram before transmission, a highly efficient use of the communications overhead is possible. This presents a very large cost savings, especially where the network is a wireless one, where it is expensive (as compared to wired networks) to transmit each datagram.

25

Brief Description of the Drawings

A preferred embodiment of the present invention will now be explained in detail with reference to the following figures, through which the invention can be better understood.

30

Figure 1 is a block diagram of a plurality of data communications devices connected to a network, and it is in this context in which the preferred embodiment of the present invention will be described;

35

Figure 2 is a flowchart illustrating the steps involved in packing data records into datagrams, according to a preferred embodiment of the present invention; and

40

Figure 3 illustrates an example of the data packing carried out according to a preferred embodiment of the present invention.

Detailed Description of the Preferred Embodiments

5 As shown in Fig. 1, a data communications apparatus 10 (which could be a personal computer) is connected to a network 11 (which could be the Internet or any other network for connecting data communications equipment). Also connected to network 11 are other data communications apparatus 12, 13 and 14. Each data communications apparatus can communicate with other data communications apparatus by sending data over the network 11.

10 Data communications apparatus 10 is shown in more detail and it will be assumed in the following discussion that this apparatus 10 is sending data to at least one other data communications apparatus (12, 13 and/or 14). Apparatus 10 includes a storage unit 101 for storing data which is to be transmitted by the apparatus 10. This storage unit 101 could be a hard disk drive, floppy disk drive, semiconductor memory, or any other well known storage unit. A data accessing unit 102 selects which data is to be transmitted and accesses such data from the storage unit 101 and then passes the accessed data along to a data record forming unit 103. Data record forming unit 103 breaks the data up into data records of various sizes, e.g., some data records will be smaller in size than others. The data record forming unit 103 then sends the data records to the data packing unit 104 which packs the data records into datagrams (e.g., unit 104 takes the data records and assigns them to the network 11's native data transmission unit, which is herein called a datagram).

15

20

25

30 The function carried out by the data packing unit 104 when the unit 104 is supplied with data records from the data record forming unit 103 will now be described in conjunction with the flowchart of Fig. 2. However, first a source code version of the functionality will be given in order to clearly illustrate such functionality:

35

```
40 WHILE (data_records_to_send) DO
    available_space = datagram_size;
    WHILE (available_space > smallest data record left) DO
        get largest data record that will fit;
        available_space = available_space - size of (data_record);
    ENDWHILE;
    send datagram;
ENDWHILE;
```

In Fig. 2, the process of packing data records into a first datagram begins. At step 200, a first datagram is prepared for packing (i.e., the data packing unit 104 concentrates on a first datagram which it will now assign data records to). It is determined at step 201 whether the available space in this first datagram (e.g., part of the datagram in which data records have not yet been written) is greater than or equal to the smallest data record that has not yet been assigned to a datagram. That is, it is determined whether there is still room to fit another data record into this first datagram. If the answer is YES, then control proceeds to step 202 where the largest unassigned data record that will fit in the available space in the first datagram is assigned to the first datagram. If the answer is NO at step 201, then control flows to step 204, where the first datagram is sent out of the data packing unit 104 and further out onto the network 11 for transmission to a receiving apparatus (e.g. 12). In this latter case, there is no room left in this datagram to pack in any more of the data records so the datagram is sent out over the network 11. After step 204 is carried out, control flows to step 205 where a second datagram is prepared for packing, and then control flows back to step 201.

Once step 202 is carried out, control flows to step 203 where it is determined whether there are any remaining data records which are still waiting to be assigned to a datagram. If YES, then control loops back to step 201. If NO, then control flows to step 206 where it is determined whether there are any datagrams which have at least one data record packed therein but which has not yet been sent out over the network 11. If YES, then such a datagram is sent out over the network at step 207 and then the flow of control ends. If NO, then the flow of control ends.

An example will now be taken up, in order to provide an even clearer understanding of the functionality involved within the data packing unit 104 of the preferred embodiment of the present invention. In the example, the data accessing unit 102 accesses, from storage unit 101, a data collection consisting of 81 bytes. These 81 bytes are to be transmitted from the data communications apparatus 10 to another data communications apparatus 13 over the network 11.

The data accessing unit 102 transfers the 81 byte data collection to the data record forming unit 103 which breaks up the collection into 8 data records DR1 through DR8 of various sizes, as shown near the top of Fig. 3. The first data record, DR1, represents the first 8 bytes of the

5

81 byte collection. The second data record, DR2, represents the next 4 bytes of the 81 byte collection. The third data record, DR3, represents the next 7 bytes of the 81 byte data collection. This pattern continues up to the last data record DR8 which represents the last 9 bytes of the 81 byte data collection.

10

15

The data record forming unit 103 then forwards these eight data records DR1 to DR8 on to the data packing unit 104 which will pack the data records DR1 to DR8 into datagrams for transmission over the network 11. For a first datagram (datagram 1, near the bottom of Fig. 3) of 35 bytes in length, control begins at step 201 where it is determined that there is plenty of available space left in datagram 1 (since there has yet been no data records packed in datagram 1) and thus the YES branch is taken. At step 202, the largest unassigned data record, e.g. DR5 of 20 bytes in length, is assigned to datagram 1. At step 203, it is determined that there are still seven other data records which need to be packed so control loops back to step 201.

20

25

30

At step 201, it is determined that the available space in datagram 1 is 15 bytes (e.g., 35 bytes minus 20 bytes equals 15 bytes) and that this 15 bytes is greater than the smallest data record available (which is DR7 having a 3 byte length). Therefore, control flows to step 202, where the largest unassigned data record that will fit in the datagram is assigned to datagram 1. In this case, since 15 bytes is available in datagram 1, DR3, having a length of 15 bytes, is the largest available datagram that will fit in datagram 1, so DR3 is assigned to datagram 1. Then control flows to step 202 where it is determined that there are still data records remaining to be assigned, so control flows back to step 201.

35

40

This time, at step 201, it is determined that the available space in datagram 1 is NOT greater than or equal to the smallest data record available (because there is no available space left in datagram 1 and 0 bytes is NOT greater than or equal to 3 bytes). Thus, control flows to step 204 where datagram 1 is sent over network 11 to data communications apparatus 13. Then, control flows to step 205 where another datagram, datagram 2 (which also has a 35 byte length) is prepared for packing. Control then returns to step 201, where it is determined whether the available space in datagram 2 is greater than the smallest data record available. In this case, the answer is YES since the whole 35 bytes of datagram 2 is available for data record packing, and thus control flows

to step 202. At step 202, the largest unassigned data record (e.g., DR6 of 15 bytes) is assigned to datagram 2. At step 203, it is then determined that there are other data records remaining to be assigned, and thus control flows back to step 201.

5

At step 201, it is determined that the available space in datagram 2, which is 20 bytes (35 bytes minus 15 bytes is 20 bytes) is greater than the smallest data record available (DR7 of 3 bytes) since 20 bytes is greater than 3 bytes. Thus, control flows to step 202 where the largest unassigned data record (DR8 of 9 bytes length) is assigned to datagram 2. At step 203, since there are still data records to be assigned, control loops back to step 201.

10

At step 201, it is determined that the available space in datagram 2, which is 11 bytes (20 bytes minus 9 bytes is 11 bytes) is greater than the smallest data record available (DR7 of 3 bytes) since 11 bytes is greater than 3 bytes. Thus, control flows to step 202 where the largest unassigned data record (DR1 of 8 bytes length) is assigned to datagram 2. At step 203, since there are still data records to be assigned, control loops back to step 201.

20

At step 201, it is determined that the available space in datagram 2, which is 3 bytes (11 bytes minus 8 bytes is 3 bytes) is equal to the smallest data record available (DR7 of 3 bytes). Thus, control flows to step 202 where the largest unassigned data record that will fit (DR7 of 3 bytes length) is assigned to datagram 2. At step 203, since there are still data records to be assigned, control loops back to step 201.

25

This time, at step 201, it is determined that the available space in datagram 2 is NOT greater than or equal to the smallest data record available (because there is no available space left in datagram 2). Thus, control flows to step 204 where datagram 2 is sent over network 11 to data communications apparatus 13. Then, control flows to step 205 where another datagram, datagram 3 (having a length of 30 bytes) is prepared for packing. Control then returns to step 201.

35

At step 201, it is determined that the available space in datagram 3 (the entire 30 bytes) is greater than the smallest data record available (which is now DR2 of 4 bytes in length) and thus control flows to step 202 where the largest unassigned data record (e.g., DR4 of 7 bytes) is assigned to datagram 3. Control then flows to step 203 where

40

it is determined that there is still one more data record (e.g., DR2) remaining to be assigned so control flows to step 201. At step 201, it is determined that the available space in datagram 3 is 23 bytes (30 bytes minus 7 bytes is 23 bytes) and this is greater than the smallest data record available (DR2 has 4 bytes) so control flows to step 202 where DR2 is assigned to datagram 3.

Now, at step 203, the NO branch is taken since there are no more data records remaining to be assigned. Control thus flows to step 206 where it is determined that there is a datagram (e.g., datagram 3) which has not yet been sent but which has been packed with at least one data record (e.g., DR4 and DR2). Thus, the YES branch is taken and control flows to step 207 where datagram 3 is sent over the network 11 to data communications apparatus 13, and finally the flow of control is ended.

The receiving data communications apparatus 13 need only remove each data record, in turn, from received datagrams and insert the data records into the correct place in a linked list of data records. The data can then be "streamed out" in the correct order in whatever sized pieces are appropriate. Accordingly, at the transmitting end, a small header should be added to each datagram and also to each data record within the datagram in order to identify the data contained. In many networks (such as the Internet) there is no guaranteed order of delivery of datagrams, and thus the receiving apparatus 13 need only check the datagram header to determine the proper order.

For large collections of data to be transmitted, the data record forming unit 103 could, preferably, group data records into bunches and pass a single bunch of data records at one time along to the data packing unit for packing into datagrams. The data packing unit would then carry out the flowchart of Fig. 2 on only the one bunch at a time. When a transmitted bunch of data records is received by a receiving data communications apparatus 13, apparatus 13 sends an acknowledgement back through the network 11 to the sending apparatus 10 to inform apparatus 10 that apparatus 13 has safely and correctly received the entire bunch of data records. In this case, a bit flag can be added to the last datagram in the bunch to indicate that it is the end of the bunch and that it should be acknowledged. This technique saves on the number of acknowledgements required when compared to a prior technique of sending a separate acknowledgement back for each datagram received.

Another embodiment of the present invention uses a similar flow chart to Fig. 2 but considers the space available in two (or more) datagrams when carrying out the packing. For example, step 201 determines whether the available space in a set of two datagrams is greater than or equal to the smallest data record available. In this way, data records larger than the datagram size can be handled and also data records (even small ones that could fit entirely in one datagram) can be made to span more than one datagram.

The data communications apparatus 10 having the described functional units has been described as the preferred embodiment. However, the invention also encompasses a data communications apparatus which does not have the storage unit, accessing unit and/or data record forming unit. For example, the data communications apparatus could receive, as an input, data records (already formed) to be transmitted. As another example, the data communications apparatus could receive, as an input, data which has already been accessed from storage but which still needs to be formed into records and packed into datagrams.

While in the preferred embodiment at step 202 the largest unassigned data record that will fit in the datagram presently being packed is selected, another embodiment could select some other criteria for determining the next candidate data record which will be checked to see if it will fit in the datagram (or data records could be taken at random and the only criteria would be whether the data record fits into the datagram).

CLAIMS

5 1. A data communications apparatus connected to a network through which datagrams flow from the data communications apparatus to at least one other data communications apparatus also connected to the network, the apparatus comprising:

10 means for receiving data which is to be transmitted over the network, such data being broken up into data records; and

15 means for packing the data records into datagrams for transmission over the network, including a means for assigning data records to datagrams;

20 wherein the means for packing takes the data records out of sequential order when assigning data records to datagrams.

25 2. The apparatus of claim 1 wherein the means for packing assigns, to a datagram presently being packed, the largest unassigned data record that will fit in the datagram.

30 3. The apparatus of claim 1 wherein the means for assigning considers only one datagram at a time when assigning data records to datagrams.

35 4. The apparatus of claim 1 wherein the means for assigning considers a plurality of datagrams at a time when assigning data records to datagrams.

40 5. The apparatus of claim 1 wherein the size of each data record must be less than or equal to the size of each datagram.

45 6. The apparatus of claim 1 wherein the size of a data record is larger than the size of a datagram.

50 7. A data communications method of operating a data communications apparatus connected to a network through which datagrams flow from the data communications apparatus to at least one other data communications apparatus also connected to the network, the method comprising steps of:

receiving data which is to be transmitted over the network, such data being broken up into data records; and

5 packing the data records into datagrams for transmission over the network, including assigning data records to datagrams;

wherein the packing step takes the data records out of sequential order when assigning data records to datagrams.

10 8. A computer program product stored on a computer readable storage medium for, when run on a computer, carrying out a data communications method of operating a data communications apparatus connected to a network through which datagrams flow from the data communications apparatus to at least one other data communications apparatus also connected to the network, the method comprising steps of:

15 receiving data which is to be transmitted over the network, such data being broken up into data records; and

20 packing the data records into datagrams for transmission over the network, including assigning data records to datagrams;

wherein the packing step takes the data records out of sequential order when assigning data records to datagrams.



Application No: GB 9814669.9
Claims searched: 1-7

Examiner: Keith Williams
Date of search: 26 January 1999

Patents Act 1977
Search Report under Section 17

Databases searched:

UK Patent Office collections, including GB, EP, WO & US patent specifications, in:

UK Cl (Ed.Q): H4P (PFD, PPEC, PPS)

Int Cl (Ed.6): H04L 12/08, 12/56

Other: Online EPODOC, WPI

Documents considered to be relevant:

Category	Identity of document and relevant passage	Relevant to claims
A	GB 2314740 A Thomson CSF - see abstract	1,7

X	Document indicating lack of novelty or inventive step	A	Document indicating technological background and/or state of the art.
Y	Document indicating lack of inventive step if combined with one or more other documents of same category.	P	Document published on or after the declared priority date but before the filing date of this invention.
&	Member of the same patent family	E	Patent document published on or after, but with priority date earlier than, the filing date of this application.