

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
12 December 2002 (12.12.2002)

PCT

(10) International Publication Number
WO 02/100064 A2

(51) International Patent Classification⁷: **H04L 29/06**

(21) International Application Number: PCT/EP02/05624

(22) International Filing Date: 22 May 2002 (22.05.2002)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
01112026.8 23 May 2001 (23.05.2001) EP

(71) Applicant (for all designated States except US): **TELEFONAKTIEBOLAGET L M ERICSSON (PUBL)**
[SE/SE]; S-126 25 Stockholm (SE).

(72) Inventors; and

(75) Inventors/Applicants (for US only): **MEYER, Michael**
[DE/DE]; Grossheidstrasse 27, 52080 Aachen (DE).

LUDWIG, Reiner [DE/DE]; Maubacher Strasse 4, 52393 Hürtgenwald (DE). **SACHS, Joachim** [DE/DE]; An den Finkenweiden 43, 52047 Aachen (DE). **SAGFORS, Mats** [FI/FI]; Ravalsvägen 8-C, FIN-02400 Kyrkslätt (FI).


(74) Agents: **HOFFMANN, Eitle** et al.; Arabellastrasse 4, 81925 München (DE).

(81) Designated States (*national*): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, OM, PH, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VN, YU, ZA, ZM, ZW.

(84) Designated States (*regional*): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW),

[Continued on next page]

(54) Title: METHOD AND SYSTEM FOR PROCESSING A DATA UNIT



(57) Abstract: A method of processing a data unit of a first protocol layer (L3) for transmission in a data unit based communication system is described, comprising the steps of passing to a second protocol layer (L2) a given data unit of said first protocol layer (L3) that is to be transmitted, said second protocol layer (L2) lying below said first protocol layer (L3), determining a numeric value of a numerically quantifiable parameter associated with said given data unit of said first protocol layer (L3), embedding said given data unit of said first protocol layer (L3) into one or more data units of said second protocol layer (L2), performing transmission control for said one or more data units of said second protocol layer (L2) that embed said given data unit of said first protocol layer (L3), where said embedding and/or said transmission control is performed in accordance with said numeric value of said numerically quantifiable parameter.

WO 02/100064 A2



Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM),
European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR,
GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent
(BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR,
NE, SN, TD, TG).

UZ, VN, YU, ZA, ZM, ZW, ARIPO patent (GH, GM, KE, LS,
MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent
(AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent
(AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU,
MC, NL, PT, SE, TR), OAPI patent (BF, BJ, CF, CG, CI,
CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG)

Declaration under Rule 4.17:

- *as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii)) for the following designations AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, OM, PH, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG,*

Published:

- *without international search report and to be republished upon receipt of that report*

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

Method and system for processing a data unit

[Field of the invention]

The present invention relates to a method of processing a data unit in a data unit based communication system, and to a corresponding data unit based communication system.

[Background of the invention]

A well known principle for data exchange in networks is that of data unit exchange. This means that data to be sent is broken down into individual units. Rules for sending and receiving such units, as well as rules for the structure of the units themselves, are determined by so-called protocols. Protocols are sets of rules that allow the communication between a sending end and a receiving end, as the rules specify how and in what form data to be sent has to be prepared such that the receiving end may interpret the data and react in accordance to protocol defined rules to which both partners in the communication adhere. The two ends of a communication adhering to a specific protocol are also referred to as peers.

Such data units are sometimes referred to by different names, depending on the type of protocol involved, such as packets, frames, segments, datagrams, etc. For the purpose of clarity the present description uses the term "data unit" generically for any type of data unit associated with any type of protocol.

An important concept in communications using data unit exchange is that of protocol layering. This means that a number of protocols (sometimes also referred to as a suite) is organised in a hierarchy of layers, where each layer has specific functions and responsibilities. The concept of

layering is well known in the art and described in many textbooks, for example "TCP-IP illustrated, volume 1, The Protocols" by W. Richard Stevens, Addison Wesley, 1994, such that a detailed description is not necessary here.

5

The TCP/IP protocol suite is an example of a layered protocol hierarchy. A basic structure of a protocol hierarchy is defined by the OSI (Open System Interconnection) layer model. At a lowest layer, which is also referred to as the physical layer or L1, the functions of directly transporting data over a physical connection are handled. Above the physical layer, a second layer L2, which is also referred to as the link layer is provided. The link layer L2 fulfils the function of handling the transport of data units over links between communication nodes. Above the link layer L2 a third layer L3 is provided, which is also referred to as the network layer. The network layer handles the routing of data units in a given network. An example of a network layer protocol is the internet protocol (IP). Above the network layer, a fourth layer L4 is provided, which is also referred to as the transport layer. Examples of a transport layer protocol are the transmission control protocol (TCP) or the user datagram protocol (UDP).

25

In a data unit based communication system using a hierarchy of protocol layers, a communication comprises passing a given data unit downwards through the protocol hierarchy on the sending side, and passing a data unit upwards through the protocol hierarchy on the receiving side. When a data unit is passed downwards, each protocol will typically perform a certain function with respect thereto, e.g. add further information and change or adapt the structure to specific rules of that protocol layer. Typically each protocol layer will add its own header to a data unit received from a higher protocol layer and may also add delimiters. When a specific protocol layer receives a data

35

- unit from a higher protocol layer, it will embed the higher layer data unit into a data unit adhering to the rules of the given protocol layer. The term "embedding" shall refer to both encapsulation in which one data unit of a higher layer is placed into one data unit of a given layer, and to segmentation, where one data unit of a higher layer is segmented into a plurality of data units of the given protocol layer.
- 10 An important aspect of the layering scheme is that the different layers are "transparent". This means that the peers in a layer are oblivious to what happens in another layer.
- 15 Typically, each protocol layer will perform some type of transmission control for its data units. Such transmission control can e.g. comprise the performing of a certain type of forward error correction, the setting of parameters associated with an automatic repeat request (ARQ) function,
- 20 the scheduling of data units, or the performing of comparable operations.

It is known to implement protocol layers in such a way that they can be operated in a specific mode with respect to the transmission control. As an example, a so-called numbered mode (or I-mode) and a so-called unnumbered mode (UI-mode) are known. In the numbered mode, if it is determined that a sent data unit was not correctly received by the receiving peer, then the sending peer performs retransmission of said data unit. In this way it can be assured that all packets are correctly transmitted, although this may increase the delay, depending on how many packets have to be transmitted. On the other hand, in the unnumbered mode, no retransmissions are provided. This has the advantage of less delay, but the transmission reliability depends on the quality of the physical connection.

The possibility of setting a given protocol layer implementation into a specific transmission control mode has the advantage that the selection of the mode can e.g. be performed by a control procedure from a higher protocol layer, in order to optimise the sending of data units from said higher protocol layer. However, it does not provide very much flexibility, as a given protocol layer will typically handle a variety of different types of data units, that require different control settings with respect to the optimisation of the sending of the given type of data unit. As an example, if an application layer is sending a computer file, it desires to ensure a reliable transmission, and may therefore want to set a lower layer protocol implementation into the numbered mode, or the application layer may want to send data that requires real-time transmission, such as a video stream belonging to a video telephone, in which case transmission speed is more important than reliability, such that the application layer may want to set a lower layer protocol implementation into the unnumbered mode. However, if both computer file data and video data are being sent, then the setting of the lower layer protocol implementation into a given transmission control mode will not provide an optimum solution.

25

EP-0 973 302 A1 addresses this problem and proposes a system, in which a given protocol layer that receives higher data unit layers and embeds these higher data layers into data units of said given layer, is arranged to discriminate the higher layer data unit layers by reading the header information and determining the type of the higher layer data unit. Then, a classification is performed in accordance with the identified type. In this way, the given protocol layer can flexibly set the transmission reliability of its data units depending on the type of the higher layer data unit that is embedded in its own data units. As an example, if the system of EP- 0 973 02 A1 is

35

applied to a link layer in the TCP/IP suite, then the link layer can identify if the network layer IP data unit that it receives carries a TCP data unit, in which case the link layer data units are sent in the numbered mode, or if the
5 network layer IP data unit carries UDP data unit, in which case the link layer data units are sent in the unnumbered mode.

However, the system of EP-0 973 302 A1 is not always
10 practical, as it requires the parsing of higher layer data units in order to identify type information, which e.g. does not work when the higher layer data unit has an encrypted header and/or payload.

15 [Object of the present invention]

It is desirable to provide an improved method and system of processing data units of a higher protocol layer at a given protocol layer, which is simple to implement, but flexibly
20 provides improved transmission properties.

[Summary of the invention]

This object is achieved by the subject-matter of the
25 independent claims. Advantageous embodiments are described in the dependent claims.

In accordance with an embodiment of the present invention, at a given protocol layer, e.g. a link layer L2, one or
30 more numeric values of one or more numerically quantifiable parameters associated with a received given data unit of a higher protocol layer, e.g. the network layer L3, are determined. In other words, one value of one numerically quantifiable parameter can be determined, or several values
35 of one numerically quantifiable parameter, or one or more respective values for each of a plurality of numerically quantifiable parameters. The at least one numeric value is

not derived from information contained in the given data unit of the higher protocol layer. In other words, instead of analysing the content of the higher layer data unit, one or more simple physical properties that are numerically
5 evaluatable are measured, and the embedding and/or transmission control is performed in accordance with the determined value. Consequently, no parsing of higher layer data units or other similar complicated processing is necessary.

10

As an example, a numerically quantifiable parameter can be the size of the higher layer data unit. In other words, at a given protocol layer e.g. L2, a higher layer data unit, e.g. from the L3 layer, is received, and the size of said
15 L3 data unit is measured. Then the embedding operation for embedding said L3 data unit into one or more L2 data units, or the transmission control operation for transmitting the one or more L2 data units into which said L3 data unit has been embedded, is performed in accordance with the result
20 of said size measurement.

Preferably, the size measurement is used as a basis for adjusting the transmission control to optimise
predetermined target properties. More specifically, the L2
25 data units are transmitted with parameters set for optimising throughput if the L3 data unit falls into a predetermined size range, and otherwise the L2 data units are transmitted with optimised delay. Namely, if the L3 data unit is found to have a size indicative of a maximum
30 size e.g. falls into a range around or is equal to the TCP maximum segment size if the L3 data units carry TCP data units, then the transmission of the L2 data units, into which said L3 data unit has been embedded, is optimised for throughput, as it may be assumed that the maximum size L3
35 data unit belongs to a larger amount of data being sent from above the L3 layer. On the other hand, if the L3 data unit is smaller in size, then the transmission control is

optimised for delay, as it can be assumed that the smaller L3 data units are associated with control operations, such as synchronisation or acknowledgment messages, where delay optimisation is more suitable than throughput optimisation.

5

Other examples of a numerically quantifiable parameter that can be used in the context of the present invention are a buffer fill level of a buffer holding data units of the upper protocol layer (e.g. L3), or of a buffer holding data units of the protocol layer (e.g. L2) receiving the upper layer data units. Another numerically quantifiable parameter is the inter-arrival time of the upper layer data units i.e. the time that passes between the arrival of two consecutive upper layer data units.

15

Each of the given examples of numerically quantifiable parameters can be used alone as a basis for providing one or more values to be used in performing the embedding and/or transmission control, or can be used together with one or more of the named numerically quantifiable parameters for providing such values to be used in performing the embedding and/or transmission control.

20

According to a preferred embodiment, the method of the present application also comprises a step of discriminating a group of data units of the higher protocol layer to which the higher layer data unit belongs. This discrimination can occur on the basis of source information and/or destination information in the higher layer data unit and/or a protocol identifier in the higher layer data unit. Especially, if the present preferred embodiment is applied in the context of TCP/IP, the discrimination can consist in determining the flow to which the higher layer data unit belongs. A flow is defined by the source and destination IP address, the source and destination port number and a protocol identifier.

30

35

The result of said discrimination can be used in the transmission control procedure for the data units into which the higher data unit layer is embedded. More specifically, referring to the above described example in which the size of the higher layer data unit was compared with a reference range or a reference value indicative of a maximum data unit size, it is possible to enhance such a comparison by taking into consideration the type of the higher layer data unit. Namely, the type of the higher layer data unit (or the type of a data unit contained in the higher layer data unit) can be used as a means for determining a reference range or reference value with which to compare the size of the higher layer data unit. As an example, if the higher layer data unit is identified as being of the first type (i.e. belonging to a first predetermined group), then the size of said higher data unit is compared to a first size reference value (or to a first set of size reference values, i.e. a discrete range of values, or to a first continuous range of values), whereas if the higher layer data unit is discriminated as being of a second type (i.e. belonging to a second predetermined group), then the size of said higher layer data unit is compared to a second size reference value (or to a second set of size reference values, or to a second continuous range of values).

On the other hand, the result of said discriminating step can also be used as a basis for determining the numerically quantifiable parameter of which a numeric value is determined. Namely, the numerically quantifiable parameter can be associated with a buffer fill level of a buffer holding data units of the lower protocol layer, where the numeric value is the number of data units of the lower protocol layer (e.g. L2) in the buffer that embed data units of the higher protocol layers belonging to the discriminated group, e.g. belonging to the discriminated flow. Equally, when using the inter-arrival time of the

higher layer data unit as the numerically quantifiable parameter, then the numeric value can be chosen as the inter-arrival time value for those higher layer data units belonging to the discriminated group. In the latter case, 5 i.e. when determining the inter-arrival time of data units belonging to a group, it is possible to add a further condition, e.g. to determine the inter-arrival time of such data units of the group that fall into a given size category. For example, if the group is a flow and the size 10 category is selected to relate to data units that contain an acknowledgement (i.e. have a minimum size), then it is possible to use the inter-arrival time of the acknowledgment data units in the flow as a basis for the transmission control.

15

The transmission control performed in accordance with the numeric value of the numerically quantifiable parameter can consist in any suitable measures such as the adjusting of forward error correction, the adjusting of ARQ settings, or 20 the adjusting of scheduling. Furthermore, if the embedding operation comprise a segmentation of the higher layer data unit into several lower layer data units, then this segmentation operation can be performed in accordance with the numeric value, namely by adjusting the size of the 25 lower layer segments in accordance with the determined value. The transmission control can consist in making adjustments for the data units at the layer receiving the higher layer data units, or at a layer below. For example, if the method of the invention is applied to a link layer 30 L2, then adjustments can be made for the L2 data units, but also for L1 data units.

As already mentioned, the transmission control can be performed to optimise certain target parameters, depending 35 on the determined numeric value. Examples of such target parameters are the throughput (the amount of payload data

transported per unit of time) or the (average) transmission delay.

According to another preferred embodiment, the transmission
5 control for transmitting the lower layer data units, into
which a higher layer data unit has been embedded, comprises
a discrimination of the lower layer data units, such that
each of the one or more lower layer data units is
classified into one of a plurality of predetermined
10 transmission categories on the basis of the discrimination
result. This discrimination of lower layer data units can
advantageously be combined with the discrimination of
higher layer data units in such a way that the lower layer
data units embedding a higher layer data unit belonging to
15 a predetermined group are themselves divided into sub-
groups associated with the group of the higher layer data
unit. For example, if the group to which the higher data
unit (e.g. L3 data unit) belongs is a flow, then such a
flow can be divided into sub-flows at the lower protocol
20 layer (e.g. L2) that embeds the data units of the flow.

According to another embodiment of the invention, one or
more numeric values of a numerically quantifiable parameter
are used in a congestion alleviation procedure, such that a
25 decision step for deciding whether to perform a congestion
alleviation measure with respect to a data unit or not
depends on said numeric value or values. The congestion
alleviation procedure can e.g. comprise a data unit
dropping procedure and/or a data unit marking procedure.
30 Therefore, the congestion alleviation measure can consist
in dropping a data unit from the buffer, or in adding a
marking to a data unit, said marking informing the
communication end-points that congestion is taking place.
An example of such a marking is the Explicit Congestion
35 Notification (ECN) known from TCP/IP.

For example, L3 data units received at layer 2 are buffered before processing. Furthermore, a congestion alleviation procedure for performing a congestion alleviation measure with respect to one or more buffered L3 data units is provided. The congestion alleviation procedure can be triggered by a number of predetermined conditions, e.g. when a link-overload of the link over which the L2 data units are to be sent is detected, or a buffer overflow of the buffer storing the received L3 data units is detected. According to the embodiment, a decision step for deciding whether to perform a congestion alleviation measure or not depends on the described numeric value, e.g. on the size of the L3 data unit, or on the inter-arrival time of L3 data units. Although the above example related to the buffering of L3 data units received at L2 (sometimes also referred to as service data units SDUs), the concept of making a congestion alleviation decision dependent on a numeric value associated with a given L3 data unit is also applicable to one or more L2 data units (L2 PDUs) in which the given L3 data unit is embedded. However, especially if the congestion alleviation procedure consists in a data unit dropping procedure, it is preferable to perform any congestion alleviation at the highest sub-layer, i.e. at the SDU level, in order to avoid unnecessary embedding operations for data that is possibly dropped.

The method and system of the present invention can be put to practice in any suitable or appropriate way by hardware, by software or by any suitable combination of hardware and software.

[Brief description of figures]

The present invention shall now be described by referring to detailed embodiments, which are only intended to be illustrative, but not limiting, with reference to the appended drawings, in which:

- Fig. 1 shows a flow chart for explaining an embodiment of the present invention;
- 5 Fig. 2 shows a flow chart for explaining another embodiment of the present invention, in which the numeric value is determined based on data unit size;
- 10 Fig. 3 shows a flow chart for explaining another embodiment of the present invention, in which higher layer data units are additionally discriminated;
- 15 Fig. 4 shows a flow chart for explaining another embodiment of the present invention, in which the numeric value is determined based on the discrimination result;
- 20 Fig. 5 shows an overview of a protocol structure to which the present invention can be applied;
- Fig. 6 schematically illustrates the concept of flow-splitting in the context of the example shown in
25 Fig. 5;
- Figs. 7a-7c
show routines of an embodiment of the invention in which data units are buffered and a data unit
30 dropping procedure is provided;
- Fig. 8 shows an example of a data unit dropping procedure;
- 35 Fig. 9 shows a further example of a data unit dropping procedure; and

Fig. 10 shows a schematic block diagram of a system that is arranged to perform the routines and procedures of Figs. 7 to 9.

5 [Detailed description of embodiments]

In the following, detailed embodiments of the present invention shall be described in order to give the skilled person a full and complete understanding. However, these
10 embodiments are illustrative and not intended to be limiting, as the scope of the invention is defined by the appended claims.

The following examples shall be described in the context of
15 TCP/IP. However, it may be noted that the present invention is applicable to any protocol hierarchy in which higher layer data units are embedded into lower layer data units, and where a transmission control is performed for the lower layer data units.

20 The following embodiments shall be described in the context of applying the invention to a link layer L2, said link layer embedding L3 data units from the network layer.

However, this is only an example, and the present invention
25 can be applied at any level of a protocol hierarchy i.e. also below or above the link layer.

Fig. 1 shows a flow chart describing a first embodiment of the present invention. As can be seen, in a first step S1 a
30 L3 data unit is passed to the link layer L2 below the layer L3. Then, in step S2, a numeric value of a numerically quantifiable parameter associated with the L3 data unit is determined. It may be noted that the term "associated with the L3 data unit" is to be understood broadly as relating
35 to any numerically quantifiable parameter derivable from the L3 data unit, where said parameter can relate to the L3 data unit as such, or to a data unit from a higher layer

than L3 embedded in the L3 data unit. Then, in step S3, the received L3 data unit is embedded into one or more L2 data units. Finally, in step S4, transmission control for the one or more L2 data units and/or L1 data units that embed
5 the received L3 data unit is performed according to the numeric value determined in step S2.

Although the example of Fig. 1 shows steps S1 to S4 in a specific sequence, this is only an example and other
10 arrangements of steps S1 to S4 are possible.

The adjustments performed in step S4 will depend on the specific circumstances and the system under consideration. For example, the transmission control may comprise
15 adjusting a forward error correction for the data units of the L2 protocol layer or for the L1 data units below the L2 protocol layer. The type of forward error correction can be chosen as is desirable or appropriate. For example, the transmission control may comprise adjusting a transmission
20 power and/or a data rate (e.g. by selecting a spreading factor in a CDMA system) over a given link and/or a degree of interleaving. The mentioned link can e.g. be a wireless link.

25 If the L2 protocol layer comprises a function of providing automatic retransmission of L2 data units under predetermined conditions, then the transmission control may comprise adjusting said retransmission function.

30 If the L2 protocol layer comprises a function for the scheduling of the L2 data units, then the transmission control may comprise adjusting said scheduling.

In the event that the embedding in step S3 is a
35 segmentation operation, then the embodiment of Fig. 1 can be arranged such that the segmentation operation is performed according to said numeric value. In other words,

the segmentation operation in dependence on the numeric value can be an alternative to making the transmission control in step S4 dependent on the numeric value, or the segmentation operation in dependence on the numeric value can be a supplement to making the transmission control in step S4 dependent on the numeric value, i.e. can be provided together with such transmission control.

The adjusting of said segmentation operation may e.g. comprise adjusting the size of the L2 data units into which a given L3 data unit is segmented.

Fig. 2 shows a flowchart of another embodiment of the present invention. The same reference signs as in Fig. 1 refer to the same steps in Fig. 2. Namely, first an L3 data unit is passed to the L2 layer in step S1. Then, in step S21, which represents an example of the more general step S2 shown in Fig. 1, the size of the received L3 data unit is measured. In other words, it is determined how much data is contained in the received L3 data unit, in any suitable dimension, such as bits or bytes.

The L3 data unit is embedded into one or more L2 units in step S3. Then, the transmission control is performed in accordance with steps S41, S42 and S43, which represent a specific example of the general procedure represented as S4 in Fig. 1. More specifically, step S41 compares the measured size of the L3 data unit with at least one reference size, where said reference size preferably represents a maximum data unit size. Then, if the size of the L3 data unit is equal to the reference size, the transmission control is optimised for throughput (step S42), and if the size of the L3 data unit is different from the reference size, then transmission control is optimised with respect to delay (step S43).

In the above example, step S41 consists in comparing the measured size with a reference value. Naturally, the measured size can also be compared to a range of values, where said range can consist of a discrete number of reference values, or can be continuous, or can be a combination of discrete and continuous values. Also, in the above example, the reference value (or range) is indicative of a maximum data unit size. However, this is only an example, and the reference value (or range) can be chosen in any suitable or desirable way, depending on the given system and circumstances. For example, the reference value can also be indicative of a minimum data unit size that identifies a data unit carrying an acknowledgment.

The optimising of transmission control for throughput can e.g. consist in reducing the amount of forward error correction and enabling ARQ for ensuring transmission reliability. On the other hand, optimising transmission control for reducing delay can consist in increasing forward error correction, but disabling ARQ. This shall be explained in more detail in the following in the specific context of TCP/IP.

Non-real-time applications often run on top of TCP, whereas real-time applications usually run on top of UDP. The main design requirement for link layer (L2) protocols for non-real time applications is throughput optimisation. This is for example achieved by tuning the trade-off between forward error correction and ARQ, so that the overall throughput is optimised. It is possible to obtain higher throughputs by weakening forward error correction and enabling ARQ. Throughput is the quantity of transmitted payload per unit of time. The throughput increases with decreasing forward error correction, as forward error correction increases overhead, which leads to a reduced amount of payload being transmitted.

In an encryption based forward error correction data is divided into blocks, where each block is the result of an encryption operation that combines the payload with redundancy information that allows a decision on the part of the receiver whether a block has been correctly received and decrypted. Then a block error rate may be defined as the rate of events in which a received block is judged as having an error. It may be noted that reduced forward error correction leads to an increased block error rate, but it can be shown that throughput can nonetheless be increased. For example, when calculating throughput as the product of data rate multiplied by $(1 - \text{block error rate})$, a data rate of 8 kbit/s and block error rate of 0.01 leads to a throughput of 7.92 kbit/s, whereas a data rate of 12 kbit/s (achieved by less forward error correction) and a block error rate of 0.1 (due to the reduced forward error correction) lead to a throughput of 10.8 kbit/s. Consequently, although the block error rate is ten times higher, the increased payload data rate leads to an overall increased throughput.

The problem is that even in a flow corresponding to a TCP based bulk data transfer, which should overall be optimised for throughput, there are individual data units or segments, which are delay sensitive. Examples are the connection set-up messages, which should be exchanged as quickly as possible to be able to start with data transmission. Another example is the TCP acknowledgement messages in reverse direction, or HTTP requests, which should be received as soon as possible.

In order to differentiate between L3 data units that should be optimised for throughput and such that should be optimised for delay reduction, the present embodiment determines the size of the L3 data unit, see step S21 in Fig. 2.

Namely, L3 data units for which the throughput should be optimised in general belong to a larger application data amount, which is segmented into transport layer data units, e.g. TCP segments. Therefore, these transport layer data units, which are then embedded in network layer (L3) data units generally have a maximum transfer unit size, e.g. the TCP maximum segments size (MSS).

Consequently, as already mentioned, step S41 of Fig. 2 can be implemented in such a way that the comparison is conducted with respect to a reference value or reference range indicative of a predetermined maximum transfer unit size. Typically, the value of the maximum transfer unit size is 256 byte, 512 byte, 536 byte or 1460 byte, such that taking into account the IP header of 40 byte, suitable reference size values used in step S41 could be 296 byte, 552 byte, 576 byte or 1500 byte.

As already mentioned, step S41 can be implemented in such a way that the measured size of the L3 data unit is compared with a plurality of reference sizes, e.g. the previously mentioned series of discrete values 296, 552, 576 and 1500 bytes, or with a suitable continuous range, or with a combination of discrete values and continuous values. Then, if the measured size of the L3 unit is equal to any one of the plurality of reference sizes, or falls into the predetermined range or ranges, then step S42 is enabled, to thereby optimise the transmission control for throughput.

In this connection, it can be mentioned that the reference size values may not only take into account the IP header, but also the possibility of header compression. In other words, the plurality of size reference values used in step S41 also contains a set of values that take into account header compression.

In the example of Fig. 2, other data units having different sizes than the reference size are sent with optimised delay (step S43). In other words, they are sent with minimum delay. In the example of TCP/IP this means that all pure
5 TCP control messages (such as synchronisation messages, acknowledgement messages, . . .) are sent with minimum delay. Also, the final data unit belonging to a larger application data amount is also transmitted with optimised (e. g. minimum) delay.

10

As already indicated above, delay optimisation can for example be achieved by choosing the appropriate coding scheme and/or adjusting the transmission power. If one assumes an additional delay for a retransmission to be 100
15 ms, one can conclude that the transmission delay becomes less and less important for small data units. Therefore, it is better to use stronger forward error correction (e.g. more redundancy) and/or higher transmission power to improve the mean delay.

20

For example, if one considers a data unit of 100 byte size, and using the previously mentioned example values of a data rate of 8 kbit/s, block error rate of 0.01 and a block size
of 20 byte, then a mean delay of $100 + 5$ ms results, or
25 assuming a data rate of 12 kbit/s, a block error rate of 0.1 and a block size of 30 byte, then a mean delay of $80 + 40$ ms results. In the first example, one needs five blocks to transmit 100 byte, in the second only 4, due to the increased data rate and therefore block size. In the first
30 case, the probability that the one of the five blocks is corrupted is approximately 5 %, leading on average to an additional delay of 5 ms. The same calculation provides an additional delay of 40 ms for the second case.
Consequently, the increased forward error correction,
35 although the data rate is reduced, leads to an improved mean delay.

Another example is when the transmission control comprises regulation via the transmission power. This can be achieved e.g. by setting a target signal-to-interference ratio for the desired link, which is maintained by a power control
5 loop. A higher signal-to-interference ratio results in an improved link quality, i.e. the number of transmission errors on the link is reduced, at the price of increased transmission power. Less transmission errors reduce the delay of error recovery and therefore the mean transmission
10 delay is reduced.

In the example of Fig. 2, the outcome of decision step S41 is to either optimise the transmission for throughput or for delay. Naturally, this is only an example. The
15 transmission control can also be optimised with regard to other target parameters, as is suitable or desired for a specific application or system. Such target parameters include system resources like e.g. transmission power, spreading codes, number of assigned physical channels.
20 Furthermore, the outcome of decision step S41 can also lead to more than the two procedures S42, S43. For example, it is possible that step S41 is implemented in such a way that the L3 data unit size is compared to n reference sizes or n reference ranges, and the outcome leads to a corresponding
25 number of n transmission control procedures, each transmission control procedure corresponding to one size reference value or range, and additionally one further default transmission control procedure for all those L3 data units that do not have the size of one of the n size
30 reference values or does not fall into one of the n ranges.

In the embodiments described above, the transmission control performed in step S4 or steps S41 to S43 was performed on the basis of the numeric value determined in
35 step S2. It should be noted that the transmission control can also be based on further measurements or determinations conducted with respect to a received L3 data unit. Namely,

as shown in figure 3, in which steps S1, S21 and S3 are identical to the corresponding steps in Fig. 2, such that a further description is not necessary, a step S5 is performed after step S3, said step S5 discriminating a group to which the L3 data unit belongs. Then, in step S40, which is an example of general step S4 of Fig. 1, the transmission control for the one or more L2 data units that embed the received L3 data unit, is performed on the basis of the numeric value determined in step S21 and the discrimination result of the discrimination step S5.

Preferably, the discrimination is performed on the basis of destination information contained in the L3 data unit and/or a protocol identifier contained in the L3 data unit. For example, the group into which the L3 data unit is classified can be the flow to which it belongs. As already mentioned previously, in the context of TCP/IP, the group into which the L3 data unit is classified can be the flow to which it belongs, and the flow is defined by a source and destination IP address, a source and destination port number, and a protocol identifier. As another example, the group to which an L3 data unit belongs can also be determined by the type of payload being transmitted, which type can e.g. be determined by checking a protocol identifier in the data unit. This means checking the protocol identifier of the L3 data unit itself, or checking one or more protocol identifiers of data units being transported as payload in the L3 data unit.

The performing of the transmission control in step S40 can be conducted in any suitable or desirable way, depending on the specific circumstances of the application. For example, procedure S40 of Fig. 3 can be arranged in a similar way as steps S41, S42 and S43 of Fig. 2, with the addition that the size reference used in the comparison is also selected in dependence on the group into which the L3 data unit is classified. For example, the group can be determined on the

basis of the type of data unit embedded in the L3 data unit, e.g. the discrimination consists in determining if the IP data unit carries a TCP data unit or a UDP data unit. If it is determined that it carries a TCP data unit, then a first set of size reference values can be used, and if it is determined that the IP data unit carries a UDP data unit, then a second set of size reference values can be used in the comparison.

10 In Fig. 3, the discrimination result obtained in step S5 is used as an element in the transmission control operation of procedure S40. However, the discrimination of a received L3 data unit into one of a predetermined number of groups can also be used as a basis for determining the numeric value that is then later to be used in the transmission control operation. This is shown in Fig. 4. After having received the L3 data unit in step S1, the discrimination step S5 is performed. Then, in step S22, which is an example of the general step S2 shown in Fig. 1, the numeric value is determined on the basis of the discrimination result. Thereafter steps S3 and S4 are conducted, as already explained in connection with Fig. 1.

For example, the numerically quantifiable parameter can be associated with a buffer fill level of a buffer holding L2 data units, where the numerically quantifiable parameter is the number of L2 data units in the buffer that embed L3 data units belonging to the group discriminated in step S5. The numerically quantifiable parameter can alternatively be the inter-arrival time of L3 data units belonging to the discriminated group. The advantages of determining the numerically quantifiable parameter in this way shall be discussed in more detail with respect to the example shown in Fig. 5.

Fig. 5 shows an embodiment of the present invention implemented in the context of a link layer operating in

accordance with the universal mobile telephone system (UMTS) as e.g. described in the Technical Specification 3G TS 25.301 V3.3.0 (1999-12), published by the 3rd Generation Partnership Project (<http://www.3gpp.org>). This document is
5 herewith fully incorporated by reference into the present disclosure.

It may be noted that Fig. 5 shows a preferred application of the present invention. However, the present invention is
10 by no means restricted thereto. Much rather, the method and system of the present invention can be applied in the context of any mobile communication device, e.g. also one operating in accordance with the general package audio service (GPRS) or any other mobile communication standard.
15 Furthermore, a method and system of the present invention are naturally not restricted to the application in mobile communication systems, but can be applied in any communication system having a protocol hierarchy.

20 In Fig. 5, a radio resource control (RRC) procedure 201 is provided, which controls the operation of other parts of the general L2 implementation, said L2 protocol implementation consisting of sub-layers L2/PDCP, L2/BMC, L2/RLC and L2/MAC. Fig. 5 schematically shows control
25 connections 204, 205, 206, 207 and 208 to the other procedures, that shall be described in the following. Namely, a one or more packet data convergence protocol (PDCP) procedures 202 are implemented for performing the functions of the packet data convergence protocol.
30 Furthermore, a BMC (Broadcast/Multicast Control) procedure 203 is implemented for performing the function of controlling the sending of data units to a plurality of destinations, via multicasting or broadcasting. Reference numeral 209 refers to radio link control (RLC) procedures
35 that implement the radio link control protocol. 211 identifies the medium access control (MAC) part of the L2

layer. Finally, 302 schematically represents the physical layer L1.

5 The radio resource control procedure 201 receives control information from the higher layer L3 (not shown). The user-plane passes down L3 data units that are to be embedded into L2 data units, to the PDCP procedure 202, to the BMC procedure 203 and/or a RLC procedure 209.

10 Reference numeral 210 represents logical channels in the L2 layer, while reference numerals 301 refers to transport channels between L2 and L1.

15 The general control for the L2 layer implementation is performed by the RRC procedure 201. More specifically, user data is transmitted on a radio access bearer. Via the radio resource control procedure 201, one or more radio bearers are configured, which determine the layer 1 and layer 2 configuration of the radio protocols. In the PDCP procedure 202, IP header compression is applied and a multiplexing of multiple traffic flows onto one logical channel is used. The RLC procedure applies backward error correction, such as ARQ, for a logical channel, among other things. The MAC procedure 211 performs a scheduling of L2 data units and
25 performs the distribution onto the transport channels 301.

As an embodiment of the present invention, the function of the PDCP procedure 202 is extended by a splitting function, which separates the received traffic flow (e. g. an IP
30 flow) into sub-flows, which are then transmitted via separate RLC connections. Since IP header compression is applied in PDCP, the necessary flow information is already available in PDCP. An example, a corresponding flow splitting is schematically depicted in Fig. 6. A flow of IP
35 data units received at the PDCP procedure 202 is split into four sub-flows, each respectively handled by a separate RLC procedure 209a, 209b, 209c and 209d.

The splitting is performed in accordance with the numeric value of the numerically quantifiable parameter, e.g. the L3 data unit size, a buffer fill level, or an inter-arrival time. Additionally, the splitting can also be conducted in dependence on the type of data unit.

For example, the numerical value can be the L3 unit size and the type of data unit being transported. Then the splitting shown in Fig. 6 can be arranged in such a way that data units having the size of an acknowledgment and being TCP data units can be transmitted with an optimum RLC configuration, i.e. minimum delay, and the other IP data units are separated onto RLC connections, which are optimised for the respective data unit size and/or respective data unit type.

An alternative way of splitting the IP flow into a plurality of RLC sub-flows is to separate different phases of the IP flow, e.g. to separate the TCP slow-start phase from the congestion avoidance phase. If the slow-start phase is configured such that a lower transmission delay is achieved, the end-to-end performance will be greatly enhanced and link resources (e.g. codes) are released more quickly.

Methods for reducing the transmission delays are e.g. a higher signal/interference-ratio target, stronger forward error correction, more aggressive ARQ settings etc. To differentiate the different phases of a TCP flow, the above described methods of employing a buffer fill-level as the numerically quantifiable parameter, or using the inter-arrival time can be used. As already mentioned previously, the determination of an inter-arrival time can also be combined with the determination of a data unit size range, in that the value used as a basis for transmission control

is the inter-arrival time of data units falling into a predetermined size range.

For example, if the RLC buffer level is used, then it could
5 be concluded that if the RLC buffer fill level is below a
predetermined threshold, the TCP flow is probably in the
slow-start phase (or at the end of application data), since
no new data is arriving. Above this predetermined
threshold, or above a second predetermined threshold, it is
10 very likely that TCP is in the congestion avoidance phase.
As a consequence, if the buffer fill level indicates a
slow-start phase, then the L2 data units are placed into an
RLC connection optimised for reducing delay, and if the
buffer fill level indicates a congestion avoidance phase,
15 then the L2 data units are placed into a RLC connection
optimised for throughput.

Corresponding considerations apply when using the inter-
arrival time of the IP data units as a numerically
20 quantifiable parameter. Namely, if the inter-arrival time
is above a predetermined threshold, then this may indicate
a slow-start phase, whereas if the inter-arrival time is
below this same predetermined threshold, or below a second
threshold, then this may indicate a congestion avoidance
25 phase.

It may be noted that an alternative to splitting the flow
into several RLC sub-flows, as shown in Fig. 6, in which
each stream of data units is handled by a separate RLC
30 procedure 209a, 209b, 209c or 209d, is to transmit all data
units in one adaptive sub-flow. In this case, all data
units are handled by the same RLC procedure, however, each
data unit is handled differently (different forward error
correction, different segmentation sizes, different ARQ,
35 different transmission power (e.g. different target signal-
to-interference ratio), etc.).

In the previous embodiments, a numeric value associated with at least one numerically quantifiable parameter associated with an L3 data unit was used as a parameter for controlling the embedding of the L3 data unit and/or
5 controlling the transmission of L2 data units embedding said L3 data unit. Now, in connection with Figures 7 to 10, embodiments will be described, which relate to applying the concept of employing a numerically quantifiable parameter to the performing of a congestion alleviation procedure for
10 buffered L3 data units and/or buffered L2 data units.

In the following examples, the congestion alleviation procedure will be described in terms of a data unit dropping procedure.

15

Figure 7a shows a routine for buffering received L3 data units. It may be noted that such received L3 (generally: higher layer) data units received at a given layer (e.g. L2) are referred to as service data units (SDUs), whereas
20 the data units into which such an SDU is embedded (encapsulated or segmented) is referred to as a protocol data unit (PDU). As a consequence, the buffer into which such received L3 data units are placed can also be referred to as an SDU buffer. This SDU buffer stores the received L3
25 data units before they are processed (embedded) into L2 data units. Returning to the routine of Fig. 7a, a first step S1 consists in passing an L3 data unit to the L2 layer, after which one or more numeric values belonging to at least one numerically quantifiable parameter associated
30 with the received L3 data unit are determined in step S2. These steps S1 and S2 are just like the steps of same reference numeral described in connection with Figures 1 to 4, and can be embodied as described in detail with respect to these Figures. Therefore, the numerically quantifiable
35 parameter can e.g. be the data unit size or the inter-arrival time of the data units. Finally, the received L3 data unit or SDU is placed into an SDU buffer in step S6,

where it can be buffered until further processing of the SDU takes place.

It may be noted that step S2 does not necessarily have to
5 be performed prior to the buffering, and could also be performed subsequently, as will be explained in connection with the embodiment of Fig. 9.

10 The buffering of received SDUs is especially advantageous in situations where the immediate processing at the receiving layer is not guaranteed or feasible, e.g. if the receiving layer is a link layer for sending data over a wireless link.

15 Fig. 7b shows a routine for processing buffered SDUs, namely for embedding the SDUs into L2 PDUs. In step S31 an SDU is taken out of the SDU buffer and embedded into one or more L2 PDUs. In step S7 the one or more resulting L2 PDUs are buffered in a PDU buffer. The buffered L2 PDUs can then
20 be transmitted according to anyone of the transmission control procedures S4, S40, S41 to S43 explained in connection with Figures 1 to 4. It may be noted that the processing routine of Fig. 7b could also be embodied by directly passing the L2 PDUs from step S31 to a
25 transmission control procedure S4, S40, S41-S43.

Fig. 7c shows a basic routine for managing the contents of one or both of the SDU buffer and PDU buffer. More specifically, in a step S8 it is determined whether a
30 triggering condition for performing dropping of buffered data units is met or not, and if it is, a data unit dropping procedure S9 is conducted. The triggering conditions can be chosen as is suitable or desirable, e.g. a data unit dropping procedure can be called for if the
35 buffer is an overflow state (e.g. the amount of data in the buffer exceeds an overflow limit), and/or if the link over which the L2 PDUs are to be transmitted is in an overload

state (e.g. the bandwidth momentarily provided by the link is lower than a predetermined limit percentage of the bandwidth demanded for sending the data that is to be transmitted over said link).

5

It may be noted that the dropping of data units is not only useful to relieve an overflow or overload state at the layer where the present embodiments are implemented, but also serves as an indication to higher layers that congestion has occurred. The slow start and congestion avoidance algorithms known from TCP/IP are examples of mechanisms with which flow control can respond to data unit loss along a transmission path.

10

15 The data unit dropping procedure S9 comprises a decision step for deciding whether a buffered data unit under consideration is to be dropped or not, where said decision step depends on one or more numeric values that belong to at least one numerically quantifiable parameter associated with a given L3 data unit. The given L3 data unit is the buffered data unit under consideration if this data unit under consideration is an L3 data unit (i.e. if the dropping of an SDU is considered, then the numerically quantifiable parameter is associated with said SDU under consideration, e.g. the size of said SDU or the inter-arrival time associated therewith), and said given L3 data unit is the L3 data unit embedded in said data unit under consideration if the data unit under consideration is an L2 data unit (i.e. if the dropping of an L2 PDU is considered, then the numerically quantifiable parameter is associated with the SDU embedded in the L2 PDU under consideration, e.g. the size of said SDU or the inter-arrival time associated therewith).

30

35 As indicated above, the data dropping procedure can be conducted in one or both of the SDU buffer and the PDU buffer. However, it is preferable to perform any data unit

dropping at the highest sub-layer, i.e. at the SDU buffer, in order to avoid unnecessary embedding operations for data that is possibly dropped.

5 Figure 8 shows an example of a data unit dropping procedure S9. In a first step S91, a first data unit is considered. The selection of the data unit to be considered can be done as is suitable or desirable, and can depend on the way the data units are buffered. For example, if the data units are
10 queued, then a data unit at a predetermined position in the queue can be selected (such as the first or last data unit), or any other type of selection routine can be chosen, such as the random selection of a data unit.

15 Then, in step S92, the above mentioned decision step is performed, i.e. it is determined whether the numeric value (e.g. data unit size or inter-arrival time) of the L3 data unit or SDU associated with the data unit under consideration (either the SDU itself or an L2 PDU embedding
20 the SDU) fulfils a predetermined condition or not. If the predetermined condition is fulfilled, the procedure branches to step S96 and drops the data unit under consideration. If not, the procedure branches to step S93, where it is determined whether there are further data units
25 present in the buffer, beyond those already considered. If yes, then the procedure branches to step S94, in which a next data unit is selected, and then the procedure loops back to step S92. The selection of a next data unit to be considered can be performed in any desirable or suitable
30 way and will generally be linked to the method used in step S91. For example, S94 can consist in simply selecting the following or preceding data unit in the queue, or in performing another random selection that only excludes previously selected data units.

35 If step S93 shows that no more data units are left that have not yet been considered for step S92 (which indicates

that none of the data units in the buffer fulfilled the dropping criterion laid out by step S92), then the procedure goes to step S95, in which a default procedure for selecting a data unit is performed, e.g. a data unit at
5 a predetermined position in a queue is selected, or a random selection is performed. Then, the thus selected data unit is dropped in step S96.

The condition checked in step S92 is preferably whether the
10 numeric value (or values) fall above or below a specific threshold. For example, if the numeric value is the size of the L3 data unit or SDU, then it is preferable to drop data units that exceed a certain size. The advantages of doing this will be explained later. If the numeric value is the
15 inter-arrival time, then it is possible to drop data units whose inter-arrival time exceeds a certain time period, or it is possible to drop data units whose inter-arrival time falls short of a certain time period.

20 In the procedure of Fig. 8, it was assumed that the numeric value is available for checking at S92. This can e.g. be the case by retrieving the value or values determined in step S2 of Fig. 7a. However, it is also possible to include
a determination step in the data dropping procedure S9.

25 This is shown in the example of Fig. 9, which relates to the case where the numeric value is the size of the L3 data unit, and the procedure comprises a step S97 for measuring the size, and where the decision step is provided in S920 as a comparison of the measured size with a predetermined
30 threshold Th, where the condition for dropping is met if the measured size exceeds the threshold Th. The remaining steps are the same as in Fig. 8 and will therefore not be described again.

35 The advantage of dropping data units that exceed a certain size (i.e. not dropping data units that fall below said certain size) is that often the dropping of small data

units will impact the data flow at higher layers much more heavily than the dropping of larger data units. Namely, small data units will often not contain payload for a data transmission (i.e. data from yet a higher layer, e.g. an application layer), but rather signalling messages that are important for setting up a connection (such as synchronisation messages), for confirming correct transmission (such as acknowledgment messages) or for releasing a connection (such as finish messages). In TCP, examples of such signalling are SYN, ACK and FIN messages. A sender of such signalling messages can often only recover from the loss of such a signalling message by performing a time-out, i.e. waiting until a predetermined time since sending the lost data unit has expired before undertaking any action to remedy the situation. As time-out periods are usually set conservatively (i.e. long), the dropping of such a signalling message can have strong impact on the overall transmission of data units, in contrast to dropping a data unit containing payload. Namely, most protocols for transporting data units are specifically equipped and designed to handle the loss of payload, such that the impact is less severe. As a consequence, overall performance can be enhanced by not dropping small data units, i.e. only dropping data units that exceed a certain size threshold.

As an example, SYN messages are sent at the beginning of a connection set up. If one is lost, the set up will not continue until a time-out has occurred. Consequently, the start of sending payload is strongly delayed. FIN messages are sent at the end of a connection, indicating a closing of the connection. Dropping such a data unit and consequently thereby providing the sender of that message with an indirect congestion indication, is of no value, as this sender is at the end of the transmission anyway. In other words, it is much better to drop a data unit from a

different source, where the congestion indication will still have an effect.

Dropping ACK messages does not improve the traffic
5 situation as TCP ACKs are cumulative.

It may be added that it is also preferable to drop larger data units in place of smaller ones because the dropping of small data units will generally not improve the buffering
10 delay at the link. In other words, more is done to alleviate the condition that triggered the data unit dropping (e.g. link overload or buffer overflow) if large data units are dropped than if small ones are dropped.

15 From the above discussion, it can be seen that the setting of the threshold Th is preferably done in dependence on the possible signalling messages sent at higher layers. As an example, if the above embodiment is applied to a layer 2 on top of which TCP is run, then the threshold Th should be
20 set such SYN, ACK and FIN messages are not dropped, i.e. the threshold Th should be set larger or equal to the expected size of the SYN, ACK and FIN messages.

It may be added that the step of discriminating a group
25 into an L3 data unit belongs (e.g. whether it belongs to a specific flow), described above in connection with step S5 in Figures 3 and 4, can also be used in connection with data unit dropping procedure. Namely, in addition to making the data unit dropping procedure dependent on the numeric
30 value, the result of such a discrimination step can also be taken into account.

For example, if the discrimination step S5 consists in determining which flow an L3 data unit belongs to, then the
35 decision step S92 or S920 can be amended to also take the discriminated flow into account, e.g. such that no data units from a specific flow or group of flows are dropped,

or that only data units from the specific flow or flow group are dropped.



Fig. 10 shows a schematic representation of a system in which the methods described in connection with Figs. 7 to 9 can be applied. The L2 implementation comprises an SDU buffer 101, an embedder 103 and a PDU buffer 104, as well as a controller 102 for controlling the operation of each of the entities 101 to 103. The controller 102 is arranged to execute suitable routines, e.g. the routine shown in Fig. 7a for buffering received SDUs in the SDU buffer 101 and the routine shown in Fig. 7b for operating the embedder 103 and PDU buffer 104. The term "implementation" refers to any hardware, software or combination of hardware and software suitable to execute the described routines and provide the desired functionality. As a consequence, the entities 101-104 can also be provided by any suitable hardware, software or combination of hardware and software.

In the above examples of Figures 7 to 10, the congestion alleviation procedure was shown as a data unit dropping procedure. Alternatively or additionally, the congestion alleviation procedure can also comprise a data unit marking procedure, in which the congestion alleviation measure does not consist in dropping the data unit, but adding an indicator or notification at a predetermined position (e.g. in a specified header field) in the data unit, where said indicator or notification informs the communication end-points of the data unit that congestion is taking place. An example of such a concept is the Explicit Congestion Notification (ECN) known in the context of TCP/IP, see e.g. RFC 3168.

Data unit marking should preferably be implemented with respect to the SDU buffer, as the SDUs will be marked if the decision step for deciding on the taking of a congestion alleviation measure decides to take a measure.

If a combination of data unit dropping and data unit marking is used a congestion alleviation procedure, then it is preferable that the SDUs be discriminated, e.g.

5 according to protocol identifiers, and SDUs belonging to a predetermined category should be dropped and not marked. In other words, data units for which it does make sense to add a marking, such as User Datagram Protocol (UDP) data units, should not be marked, but rather dropped if it is decided
10 to take a congestion alleviation measure. Namely, as is well known, UDP peers are not responsive to congestion, such that marking with a congestion notification makes no sense. With respect to data units discriminated into a category that is responsive to congestion, such as
15 Transmission Control Protocol (TCP) data units, the congestion alleviation measure can be either dropping or marking, depending on the specific preferences or desires. For example, for congestion responsive SDUs the decision step can be implemented in such a way that data unit
20 marking is performed for a first triggering condition (e.g. a triggering condition indicating that the SDU buffer is in danger of becoming congested, like the exceeding of a first threshold buffer fill level) and data unit dropping is performed for a second triggering condition (e.g.
25 exceeding a second threshold higher than the first threshold, which means that the buffer is congested, such that the buffer load must be reduced).

Although the present invention has been described with
30 reference to specific embodiments these embodiments only serve to provide the skilled person with a full and complete understanding of the invention, but are not intended to limit the scope. The scope of the invention is much rather defined by the appended claims. Reference
35 numerals in the claims are intended to make the claims easier to understand and also do not restrict the scope.

Claims

1. A method of processing a data unit of a first protocol layer (L3) for transmission in a data unit based communication system, comprising the steps of:
- 5 passing to a second protocol layer (L2) a given data unit of said first protocol layer (L3) that is to be transmitted, said second protocol layer (L2) lying
- 10 below said first protocol layer (L3);
- determining one or more numeric values, said one or more numeric values belonging to at least one numerically quantifiable parameter associated with
- 15 said given data unit of said first protocol layer (L3);
- embedding said given data unit of said first protocol layer (L3) into one or more data units of said second protocol layer (L2),
- 20 performing transmission control for said one or more data units of said second protocol layer (L2) that embed said given data unit of said first protocol
- 25 layer (L3),
- where said embedding and/or said transmission control is performed in accordance with said one or more numeric values of said at least one numerically
- 30 quantifiable parameter.
2. The method of claim 1, wherein said at least one numerically quantifiable parameter is the size of said given data unit of said first protocol layer (L3).
- 35 3. The method of claim 1 or 2, wherein said at least one numerically quantifiable parameter is associated with

a buffer fill level of a buffer holding data units of said first protocol layer (L3) or said second protocol layer (L2).

- 5 4. The method of one of claims 1 to 3, wherein said at least one numerically quantifiable parameter is an inter-arrival time of data units of said first protocol layer (L3).
- 10 5. The method of one of claims 1 to 4, furthermore comprising a step of discriminating a group of data units of said first protocol layer (L3) to which said given data unit of said first protocol layer (L3) belongs.
- 15 6. The method of claim 5, wherein said group is discriminated on the basis of source information and/or destination information and/or a protocol identifier contained in said data units of said first protocol layer (L3).
- 20 7. The method of claim 5 or 6, wherein said transmission control for said one or more data units of said second protocol layer (L2) that embed said given data unit of said first protocol layer (L3) is also performed in accordance with a result of said discriminating step.
- 25 8. The method of claim 5 or 6, wherein said numeric value is determined on the basis of a result of said discriminating step.
- 30 9. The method of claim 8, wherein said numerically quantifiable parameter is associated with a buffer fill level of a buffer holding data units of said second protocol layer (L2), said numerically quantifiable parameter being the number of data units of said second protocol layer (L2) in said buffer that
- 35

embed data units of said first protocol layer (L3) belonging to said group.

- 5 10. The method of claim 8, wherein said numerically quantifiable parameter is the inter-arrival time of data units of said first protocol layer (L3) belonging to said group.
- 10 11. The method of one of claims 1 to 10, wherein said transmission control comprises adjusting a forward error correction for said data units of said second protocol layer (L2) or for data units of a third protocol layer (L1) below said second protocol layer (L2).
- 15 12. The method of claim 11, wherein a function (RLC) controls the sending of said data units of said second protocol layer (L2) over a link, and said transmission control comprises adjusting a transmission power and/or a data rate over said wireless link and/or a degree of interleaving.
- 20 13. The method of one of claims 1 to 12, wherein said second protocol layer (L2) comprises a function of providing automatic retransmission of data units of said second protocol layer (L2) under predetermined conditions, and where said transmission control comprises adjusting said retransmission function.
- 25 14. The method of one of claims 1 to 13, wherein said second protocol layer (L2) comprises a function (MAC) for scheduling of said data units of said second protocol layer (L2), and where said transmission control comprises adjusting said scheduling.
- 30 15. The method of one of claims 1 to 14, wherein said embedding comprises a segmentation operation for data
- 35

units of said first protocol layer (L3), and where said transmission control comprises adjusting said segmentation operation.

- 5 16. The method of claim 15, wherein the adjusting of said segmentation operation comprises adjusting the size of the data units of said second protocol layer (L2) into which said given data unit of said first protocol layer (L3) is segmented.

10

17. The method of one of claims 1 to 16, wherein said transmission control comprises discriminating said one or more data units of said second protocol layer (L2) in which said given data unit of said first protocol layer (L3) is embedded on the basis of said numeric value, and placing each of said one or more data units of said second protocol layer (L2) into one of a plurality of predetermined transmission categories on the basis of said discrimination result.

20

18. A computer program arranged to execute the method of one of claims 1 to 17.

25

19. A computer program storage medium storing a computer program according to claim 18.

30

35

20. A data unit based communication system comprising implementations of a first protocol layer (L3) and a second protocol layer (L2), said first protocol layer (L3) lying above said second protocol layer (L2), said system being arranged to pass a given data unit of said first protocol layer (L3) that is to be transmitted to said second protocol layer (L2), and said implementation of said second protocol layer (L2) being arranged to determine one or more numeric values belonging to at least one numerically quantifiable parameter associated with said given data unit of said

first protocol layer (L3), embed said given data unit of said first protocol layer (L3) into one or more data units of said second protocol layer (L2), and perform transmission control for said one or more data units of said second protocol layer (L2) that embed said given data unit of said first protocol layer (L3) in accordance with said one or more numeric values of said at least one numerically quantifiable parameter.

21. The data unit based communication system of claim 20, wherein said at least one numerically quantifiable parameter is the size of said given data unit of said first protocol layer (L3).

22. The data unit based communication system of claim 20 or 21, wherein said at least one numerically quantifiable parameter is associated with a buffer fill level of a buffer holding data units of said first protocol layer (L3) or said second protocol layer (L2).

23. The data unit based communication system of one of claims 20 to 22, wherein said at least one numerically quantifiable parameter is an inter-arrival time of data units of said first protocol layer (L3).

24. The data unit based communication system of one of claims 20 to 23, wherein said implementation of said second protocol layer (L2) is arranged to discriminate a group of data units of said first protocol layer (L3) to which said given data unit of said first protocol layer (L3) belongs.

25. The data unit based communication system of claim 24, wherein said implementation of said second protocol layer (L2) is arranged to discriminate said group on the basis of source information and/or destination

information and/or a protocol identifier contained in said data units of said first protocol layer (L3).

26. The data unit based communication system of claim 24 or 25, wherein said implementation of said second protocol layer (L2) is arranged to also perform said transmission control for said one or more data units of said second protocol layer (L2) that embed said given data unit of said first protocol layer (L3) in accordance with a result of said discriminating.
27. The data unit based communication system of claim 24 or 25, wherein said implementation of said second protocol layer (L2) is arranged to determine said numeric value on the basis of a result of said discriminating step.
28. The data unit based communication system of claim 27, comprising a buffer for holding data units of said second protocol layer (L2), said numerically quantifiable parameter being the number of data units of said second protocol layer (L2) in said buffer that embed data units of said first protocol layer (L3) belonging to said group.
29. The data unit based communication system of claim 27, comprising a timer for measuring an inter-arrival time of data units of said first protocol layer (L3) belonging to said group, where said numerically quantifiable parameter is the inter-arrival time of data units of said first protocol layer (L3) belonging to said group.
30. The data unit based communication system of one of claims 20 to 29, wherein said transmission control comprises adjusting a forward error correction for said data units of said second protocol layer (L2) or

for data units of a third protocol layer (L1) below said second protocol layer (L2).

- 5 31. The data unit based communication system of claim 30, further comprising a function (RLC) for controlling the sending of said data units of said second protocol layer (L2) over a link, and said transmission control comprises adjusting a transmission power and/or a data rate over said link and/or a degree of interleaving.
- 10 32. The data unit based communication system of one of claims 20 to 31, wherein said implementation of said second protocol layer (L2) comprises a function of providing automatic retransmission of data units of
15 said second protocol layer (L2) under predetermined conditions, and where said transmission control comprises adjusting said retransmission function.
- 20 33. The data unit based communication system of one of claims 20 to 32, wherein said implementation of said second protocol layer (L2) comprises a function (MAC) for scheduling of said data units of said second protocol layer (L2), and where said transmission
25 control comprises adjusting said scheduling.
- 30 34. The data unit based communication system of one of claims 20 to 33, wherein said implementation of said second protocol layer (L2) is arranged to perform a segmentation operation for data units of said first protocol layer (L3), and where said transmission control comprises adjusting said segmentation
35 operation.
- 35 35. The data unit based communication system of claim 34, wherein the adjusting of said segmentation operation comprises adjusting the size of the data units of said

second protocol layer (L2) into which said given data unit of said first protocol layer (L3) is segmented.

5 36. The data unit based communication system of one of claims 20 to 35, wherein said transmission control comprises discriminating said one or more data units of said second protocol layer (L2) in which said given data unit of said first protocol layer (L3) is embedded on the basis of said numeric value, and
10 placing each of said one or more data units of said second protocol layer (L2) into one of a plurality of predetermined transmission categories on the basis of said discrimination result.

15 37. The data unit based communication system of one of claims 20 to 35, wherein said system is contained in a mobile communication device.

20 38. The data unit based communication system of claim 37, wherein said mobile communication device operates in accordance the standard of the Universal Mobile Telecommunication System.

25 39. A method of processing a data unit of a first protocol layer (L3) for transmission in a data unit based communication system, comprising the steps of:

30 passing to a second protocol layer (L2) a given data unit of said first protocol layer (L3) that is to be transmitted, said second protocol layer (L2) lying below said first protocol layer (L3);

performing a buffering operation (S6, S7) for one or both of


35 - said data unit of said first protocol layer (L3) and

- one or more data units of said second protocol layer into which said data unit of said second protocol layer (L2) is embedded,

5 performing a congestion alleviation procedure (S9) for buffered data units if one or more predetermined triggering conditions are fulfilled, said congestion alleviation procedure (S9) comprising a decision step (S92) for deciding whether congestion alleviation
10 measure is to be performed with respect to a buffered data unit under consideration or not, where said decision step depends on one or more numeric values that belong to at least one numerically quantifiable parameter associated with a given data unit of said
15 first protocol layer (L3), said given data unit of said first protocol layer (L3) being the data unit under consideration if the data unit under consideration is a data unit of the first protocol layer (L3), and said given data unit of said first
20 protocol layer (L3) being a data unit of said first protocol layer (L3) embedded in said data unit under consideration if the data unit under consideration is a data unit of said second protocol layer (L2).

25 40. The method of claim 39, wherein said congestion alleviation procedure comprises a data unit dropping procedure for dropping said buffered data unit under consideration and/or a data unit marking procedure for adding a marking to said given data unit of said first
30 protocol layer (L3).

41. The method of claim 39 or 40, wherein said at least
35 one numerically quantifiable parameter is the size of said given data unit of said first protocol layer (L3).




42. The method of claim 41, wherein said decision step (S92) comprises comparing (S920) the size of said given data unit of said first protocol layer (L3) with a predetermined threshold (Th), and performing the congestion alleviation measure with respect to the data unit under consideration if the threshold (Th) is exceeded.
43. The method of one of claims 39 to 42, wherein said at least one numerically quantifiable parameter is an inter-arrival time of data units of said first protocol layer (L3).
44. The method of one of claims 39 to 43, furthermore comprising a step of discriminating a group of data units of said first protocol layer (L3) to which said given data unit of said first protocol layer (L3) belongs.
45. The method of claim 44, wherein said group is discriminated on the basis of source information and/or destination information and/or a protocol identifier contained in said data units of said first protocol layer (L3).
46. The method of claim 44 or 45, wherein said decision step of said congestion alleviation procedure also depends on a result of said discriminating step.
47. The method of claim 46, wherein said group is discriminated on the basis of a protocol identifier in said given data unit of said first protocol layer (L3), and said decision step comprises determining whether the congestion alleviation measure should consist in dropping the data unit under consideration or marking the given data unit of the first protocol layer (L3), where the data unit under consideration is

dropped if the protocol identifier belongs to a predetermined category.

48. A data unit based communication system comprising
5 implementations of a first protocol layer (L3) and a second protocol layer (L2), said first protocol layer (L3) lying above said second protocol layer (L2), said system being arranged to pass a given data unit of said first protocol layer (L3) that is to be
10 transmitted to said second protocol layer (L2), and said implementation of said second protocol layer (L2) comprising a buffer (101, 104) for buffering one or both of
- said data unit of said first protocol layer
15 (L3) and
 - one or more data units of said second protocol layer into which said data unit of said second protocol layer (L2) is embedded,
- 20 a controller (102) connected to said buffer (101, 104) for performing a congestion alleviation procedure (S9) for buffered data units if one or more predetermined triggering conditions are fulfilled, said congestion
alleviation procedure (S9) comprising a decision step
25 (S92) for deciding whether a congestion alleviation measure is to be performed with respect to a buffered data unit under consideration or not, where said decision step depends on one or more numeric values that belong to at least one numerically quantifiable
30 parameter associated with a given data unit of said first protocol layer (L3), said given data unit of said first protocol layer (L3) being the data unit under consideration if the data unit under consideration is a data unit of the first protocol
35 layer (L3), and said given data unit of said first protocol layer (L3) being a data unit of said first protocol layer (L3) embedded in said data unit under

consideration if the data unit under consideration is
a data unit of said second protocol layer (L2).

1/11

**Fig. 1**

2/11




Fig. 2

3/11






Fig. 3

4/11

**Fig. 4**

5. 6. 7.



6/11







Fig. 6

7/11

**Fig. 7a****Fig. 7b**

8/11

**Fig. 7c**

9/11




Fig. 8

10/11




Fig. 9

11/11




Fig. 10