



US 20130279341A1

(19) **United States**

(12) **Patent Application Publication**

Bahr et al.

(10) **Pub. No.: US 2013/0279341 A1**

(43) **Pub. Date: Oct. 24, 2013**

(54) **CONGESTION NOTIFICATION ELEMENT AND METHOD FOR CONGESTION CONTROL**

(76) Inventors: **Michael Bahr**, Munchen (DE); **Barbara Staehle**, Furstenfeldbruck (DE); **Dirk Staehle**, Furstenfeldbruck (DE)

(21) Appl. No.: **13/991,479**

(22) PCT Filed: **Dec. 6, 2011**

(86) PCT No.: **PCT/EP11/71950**

§ 371 (c)(1),
(2), (4) Date: **Jul. 10, 2013**

(30) **Foreign Application Priority Data**

Dec. 7, 2010 (EP) 10194018.7

Publication Classification

(51) **Int. Cl.**
H04W 28/02 (2006.01)
(52) **U.S. Cl.**
CPC *H04W 28/0289* (2013.01)
USPC **370/237**

(57) **ABSTRACT**

A congestion notification element for indicating a congestion status of a mesh node may include at least one field specifying a mesh destination for which an intra-mesh congestion control is to be applied. This may allow for forwarding data frames on paths that share some links but not the bottleneck link with a path that is congestion controlled, which has not been possible with currently specified congestion notification techniques.

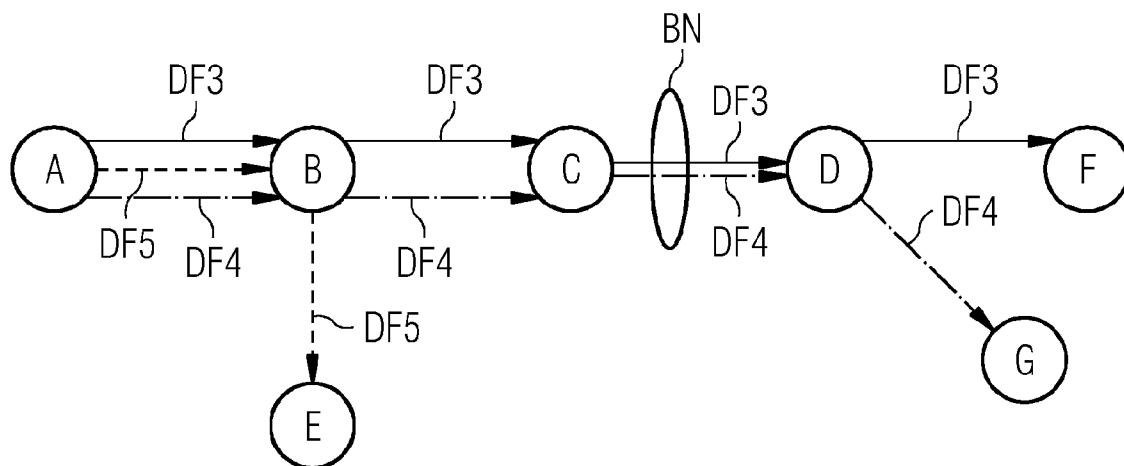


FIG 1

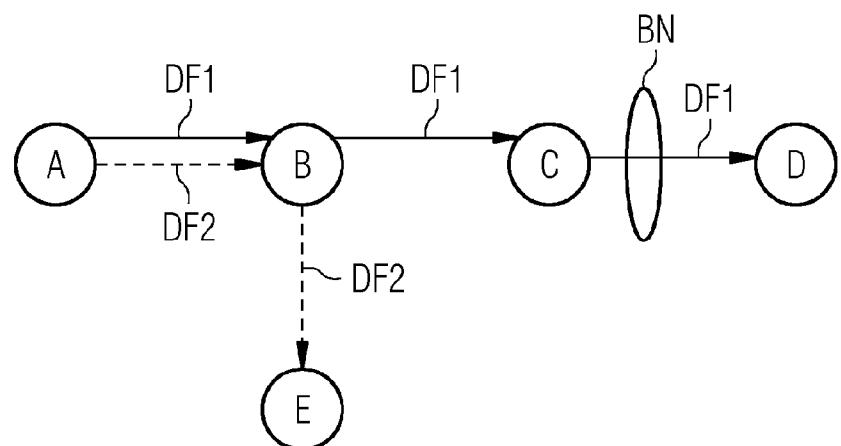


FIG 2

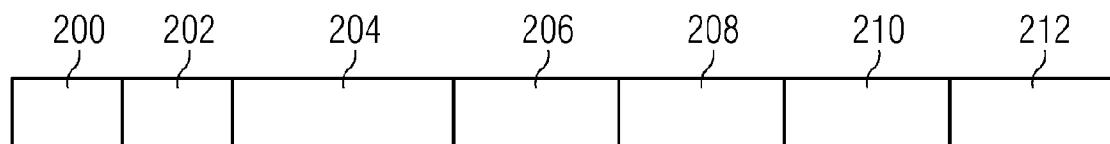


FIG 3

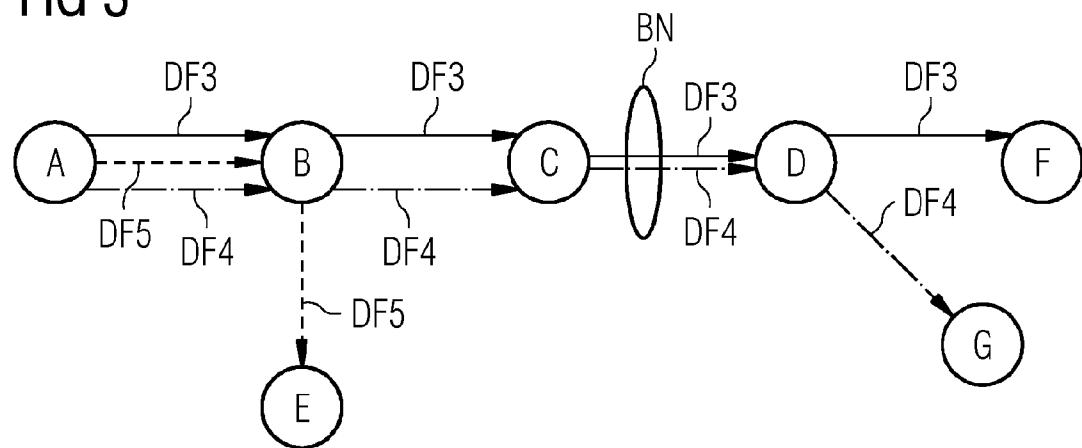


FIG 4

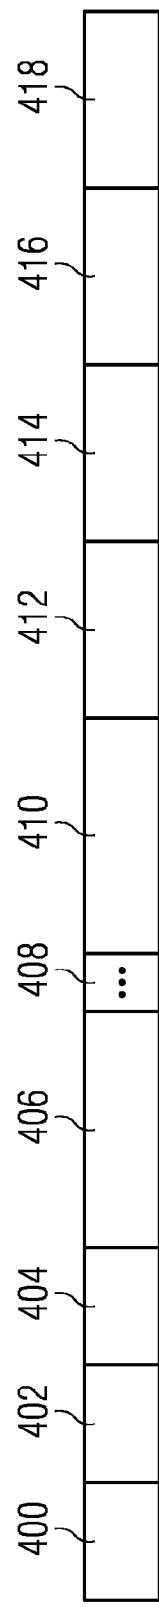


FIG 5

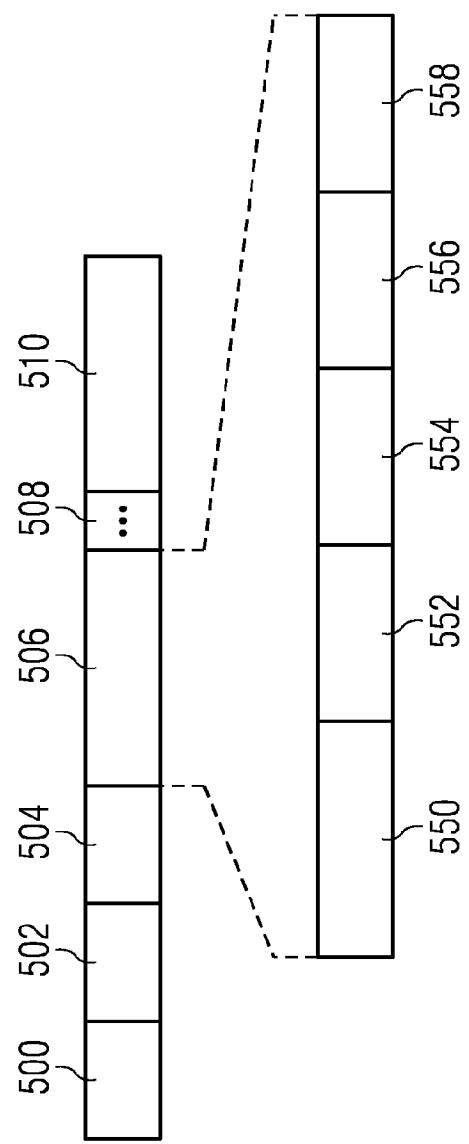


FIG 6

Element ID	14	mesh STA F	<t0> μ s	<t1> μ s	<t2> μ s	<t3> μ s
600	602	604	606	608	610	612

FIG 7

Element ID	14	mesh STA G	<t4> μ s	<t5> μ s	<t6> μ s	<t7> μ s
700	702	704	706	708	710	712

CONGESTION NOTIFICATION ELEMENT AND METHOD FOR CONGESTION CONTROL

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application is a U.S. National Stage Application of International Application No. PCT/EP2011/071950 filed Dec. 6, 2011, which designates the United States of America, and claims priority to EP Patent Application No. 10194018.7 filed Dec. 7, 2010. The contents of which are hereby incorporated by reference in their entirety.

TECHNICAL FIELD

[0002] This disclosure relates to a method for congestion control for the avoidance of traffic congestion within a network, particularly a wireless mesh network. The invention further relates to a congestion notification element granting an enhanced congestion control.

BACKGROUND

[0003] A number of congestion control methods are known in the art. An exemplary congestion control method including a congestion notification element is disclosed in IEEE P802.11s, Draft Standard >>Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications—Mesh Networking<<, version D7.03, November 2010, hereinafter referred to as >>draft standard<<. The structure of said congestion notification element is particularly described in section 7.3.2.99 of the draft standard.

[0004] In accordance with the aforementioned congestion control method, a mesh node that detects congestion may transmit a congestion notification element to the mesh nodes of its traffic source or other adjacent mesh nodes.

[0005] A major drawback of this known congestion control method is an inherent imprecision in decreasing the congestion between the source and the mesh node detecting the congestion while affecting traffic between adjacent nodes which previously have not been affected by the congestion.

[0006] Reference will now be made to FIG. 1 which demonstrates the drawback mentioned above when applying said known congestion control method.

[0007] According to FIG. 1 a plurality of mesh nodes A, B, C, D, E are logically interconnected by links which are depicted by lines interconnecting some of the mesh nodes A, B, C, D, E. A first dataflow DF1 is ranging from a first mesh node A passing a second mesh node B and a third mesh node C to a fourth mesh node D, which is the destination node D of said first dataflow DF1. A second dataflow DF2, illustrated by a dotted line, is ranging from the first mesh node A passing the second mesh node B to a fifth mesh node E, which is the destination node E of said second dataflow DF2.

[0008] A link between the third mesh node C and the fourth mesh node D is affected by a reduced transfer rate because of a bottleneck BN of any kind between the third mesh node C and the fourth mesh node D. This bottleneck BN causes congestion at the third mesh node C. In order to decrease this congestion, the third mesh node C sends a—not shown—congestion notification to the source of its link, which is the adjacent second mesh node B.

[0009] Subsequently, the second mesh node B stops or decreases, or, in other words, postpones, sending data frames to the third mesh node C. These are data frames of the first

data flow DF1. However, data frames are still sent by the second mesh node B to a fifth mesh node E along the second dataflow DF2 since the link between the second mesh node B and the fifth mesh node E is not affected by the bottleneck BN between the third mesh node C and the fourth mesh node D.

[0010] Inevitably, the second mesh node B is affected by congestion since the second mesh node B does not forward data frames to the third mesh node C any more but still receives data frames from an adjacent first mesh node A. Subsequently, the second mesh node B sends a congestion notification to the first mesh node A. Having received this notification, the first mesh node A will stop sending data frames to the second mesh node B. These are data frames of both the first data flow DF1 and the second data flow DF2.

[0011] The described scenario eventually leads to a situation, where the second data flow DF2 between the first mesh node A and the fifth mesh node E is stopped although the initial bottleneck between the third mesh node C and the fourth mesh node D has no negative impact on the second data flow DF2.

SUMMARY

[0012] One embodiment provides a congestion notification element for indicating a congestion status of a mesh node in a mesh network, the congestion notification element including at least one field specifying a mesh destination for which an intra-mesh congestion control is to be applied.

[0013] In a further embodiment, the mesh destination is determined by an address of a destination mesh node.

[0014] In a further embodiment, the mesh destination is determined by a broadcast address.

[0015] In a further embodiment, the congestion notification element includes a field specifying a number of fields each field of the number of fields specifying a respective mesh destination.

[0016] In a further embodiment, the congestion notification element includes at least one congestion notification, each congestion notification including one of the at least one field specifying the mesh destination, each congestion notification further including at least one field specifying a congestion notification expiration timer, the congestion notification element further including a field specifying a number of congestion notifications included in the congestion notification element.

[0017] Another embodiment provides a congestion control notification frame including at least one congestion notification element according to one of the aforementioned claims.

[0018] In a further embodiment, the congestion control notification frame includes fields according to known document IEEE P802.11s.

[0019] Another embodiment provides a method of controlling congestion in a mesh network, the method including the steps of: receiving at least one congestion control notification frame by a receiving mesh node, the at least one congestion control notification frame being sent by a transmitting mesh node, the at least one congestion control notification frame including at least one congestion notification element; reading at least one transmitter mesh node address of said at least one congestion control notification frame as next hop address; reading at least one destination mesh node address of said at least one congestion notification element, said destination mesh node address being an individual address or a

broadcast address; and reading at least one congestion notification expiration timer of said at least one congestion notification element.

[0020] In a further embodiment, the method includes the steps of: receiving a data frame by said receiving mesh node, the data frame including a destination address of said data frame; reading the destination address of said data frame; and reading a next hop address for the destination address of said data frame from a forwarding information.

[0021] In a further embodiment, the method includes the steps of postponing a forwarding of said data frame by said receiving node for the case that said next hop address for said data frame is equal to at least one next hop address from said at least one congestion control notification frame; and said destination address of said data frame is equal to at least one destination mesh node address from said at least one next hop address of said at least one congestion notification element.

[0022] In a further embodiment, the method includes the steps of postponing a forwarding of said data frame by said receiving node for the case that said next hop address for said data frame is equal to at least one next hop address from said at least one congestion control notification frame; and at least one destination mesh node address from said at least one next hop address of said at least one congestion notification element is equal to said broadcast address.

[0023] In a further embodiment, said postponing of said forwarding of said data frame is done until a corresponding congestion notification expiration timer of a corresponding access category has expired.

[0024] Another embodiment provides a mesh node in a mesh network comprising means for carrying out any of the methods disclosed above.

[0025] Another embodiment provides a mesh node in a mesh network comprising means for processing and generating congestion notification elements and/or congestion control notification frames as disclosed above.

[0026] Another embodiment provides a mesh network comprising at least one such mesh node.

[0027] Another embodiment provides a computer program product, which contains a program code stored on a non-transitory computer-readable medium and which, when executed on a processor of a node in a mesh network, carries out any of the methods disclosed above.

[0028] Another embodiment provides a computer program product, which contains a program code stored on a non-transitory computer-readable medium and which, when executed on a processor of a node in a mesh network, processes and/or generates congestion notification elements and/or congestion control notification frames as disclosed above.

BRIEF DESCRIPTION OF THE DRAWINGS

[0029] Exemplary embodiments will be explained in more detail below on the basis of the schematic drawings, wherein:

[0030] FIG. 1 shows an exemplary section of a network structure for illustrating a first congestion situation;

[0031] FIG. 2 shows an exemplary structure of a congestion notification element according to an example embodiment;

[0032] FIG. 3 shows an exemplary section of a network structure for illustrating a second congestion situation;

[0033] FIG. 4 shows an exemplary structure of a congestion notification element with multiple destination addresses according to an example embodiment;

[0034] FIG. 5 shows an exemplary structure of a congestion notification element with multiple congestion notifications according to an example embodiment;

[0035] FIG. 6 shows exemplary contents of a congestion notification element according to a first application example; and

[0036] FIG. 7 shows exemplary contents of a congestion notification element according to a second application example.

DETAILED DESCRIPTION

[0037] Some embodiments provide a congestion notification element for indicating an enhanced congestion status by overcoming draw-backs of currently known indications of congestions which are based on a sole next hop indication.

[0038] Other embodiments provide an improved congestion control method, which substantially does not affect links to un congested adjacent nodes.

[0039] Accordingly, an improved congestion notification element for indicating a congestion status of a mesh node in a mesh network is provided, the congestion notification element including at least one field specifying a mesh destination for which an intra-mesh congestion control is to be applied.

[0040] Accordingly, an improved congestion control method in a sense of a per-destination congestion control in mesh networks is provided, thereby overcoming the draw-backs of currently known congestion control methods based on a sole next hop control.

[0041] One possible advantage of the proposed congestion notification element and the proposed method of congestion control lies in the ability of forwarding data frames on paths that share some links but not the bottleneck link with a path that is congestion controlled by specifying a mesh destination for which an intra-mesh congestion control is to be applied. This has not been possible with the currently specified congestion notification of draft standard IEEE 802.11s.

[0042] Some embodiments may avoid negative performance impacts on flows which are not passing a congested bottleneck link. The congestion control proposed herein may be directed to data flows passing the bottleneck link.

[0043] Reference will now be made in detail to example embodiments of the present invention, examples of which are illustrated in the accompanying drawing, wherein like reference signs refer to like elements throughout.

[0044] FIG. 2 shows an exemplary structure of a congestion notification element according to an example embodiment.

[0045] The congestion notification element includes a first field 200 captioned >>Element ID<< and having a length of one octet. A second field 202 is captioned >>Length<< and has a length of one octet. The first field 200 and the second field 202 form the header of the congestion notification element. The header is followed by a payload portion of the congestion notification element, which fields are described in the following.

[0046] The congestion notification element contains four Congestion Notification Expiration Timer fields 206, 208, 210, 212, namely a first Congestion Notification Expiration Timer field 206 captioned >>Congestion Notification Expiration Timer AC_BK<<, a second Congestion Notification Expiration Timer field 208 captioned >>Congestion Notification Expiration Timer AC_BE<<, a third Congestion Notification Expiration Timer field 210 captioned >>Congestion Notification Expiration Timer AC_VI<<, and a fourth Con-

gestion Notification Expiration Timer field **212** captioned >>Congestion Notification Expiration Timer AC_VO<<.

[0047] The four Congestion Notification Expiration Timer fields **206**, **208**, **210**, **212**, are dedicated to four respective access categories to indicate the estimated congestion duration per access category at the mesh node transmitting the congestion notification.

[0048] An access category, also referred to as AC, is applied for implementing a Quality of Service Architecture (QoS). To this end a method called >>Enhanced Distributed Channel Access<<, also referred to as EDCA, is implemented to provide prioritized services using traffic category classification and prioritized media access. An access category is assigned to a respective traffic type as shown below:

Traffic Type: Background

Access Category: AC_BK

Traffic Type: Best-effort

Access Category: AC_BE

Traffic Type: Video

Access Category: AC_VI

Traffic Type: Voice

Access Category: AC_VO

[0049] The values in the Congestion Notification Expiration Timer fields **206**, **208**, **210**, **212**, are encoded as unsigned integers in units of 100 μ s.

[0050] A third field **204** is captioned >>Destination Mesh STA Address<< and has a length of six octets.

[0051] The structure of the Congestion Notification element is partially known so far from the draft standard IEEE 802.11s. According to an example embodiment, this Congestion Notification element is extended by a third field **204**, captioned >>Destination Mesh STA Address<< and having a length of six octets. This destination field **204** contains a hardware or MAC address (Media Access Control) of the mesh destination.

[0052] FIG. 2 shows a possible format of a congestion control element amended according to an example embodiment. The amended field **204** captioned >>Destination Mesh STA Address<< according to said embodiment is positioned at the first field of the payload portion of the Congestion Notification element. Of course, the amended destination field **204** can alternatively be located after the four Congestion Notification Expiration Timer fields **206**, **208**, **210**, **212** keeping the position of the already existing fields **206**, **208**, **210**, **212**.

[0053] The amended destination field **204** captioned >>Destination Mesh STA Address<< preferably has a length of six octets. Hence, the length field **202** of the Congestion Notification element is set to 14. The field **200** captioned >>Element ID<< is set to an identifier value defined for this element. The Congestion Notification element is included in—not shown—Congestion Control Notification frames.

[0054] The destination field **204** is represented as a 48-bit MAC address and is set to the address of the mesh destination for which the intra-mesh congestion control shall be applied. It is set to the broadcast address if the intra-mesh congestion control shall be applied to all destinations that use the transmitter of the congestion notification element as next hop. This broadcast address is also referred to as group address.

[0055] A mesh station receiving an amended Congestion Notification element according to an example embodiment is advantageously capable of stopping the forwarding of data frames with a better selectivity than known before. For a better understanding, a method implying the known Congestion Notification element is compared to a method implying the amended Congestion Notification element according to an example embodiment.

[0056] According to a method implying the known Congestion Notification element, a rule for postponing of data frame forwarding is described in pseudo code as follows:

```
IF (MAC address of next hop of data frame == MAC address of
transmitter of Congestion Notification element)
THEN
    do not forward data frame until the Congestion
    Notification Expiration Timer of the corresponding access
    category has expired.
ENDIF
```

[0057] The rule illustrated above will also postpone the forwarding of data frames between mesh nodes which are not affected by a bottleneck. More specifically, with reference to FIG. 1, this rule would also postpone the second dataflow DF2 between the first mesh node A and the fifth mesh node E although the initial bottleneck BN between the third mesh node C and the fourth mesh node D has no negative impact on this second dataflow DF2.

[0058] In other words, traffic between the first mesh node A and the fifth mesh node E is postponed because this traffic had to pass the second mesh node B first. However, this kind of traffic is not forwarded, because the second mesh node B has sent a Congestion Notification element to the first mesh node A.

[0059] According to a method implying the amended Congestion Notification element according to an example embodiment, an amended rule for postponing of data frame forwarding is described in pseudo code as follows:

```
IF (MAC address of next hop of data frame == MAC address of
transmitter of Congestion Notification element)
THEN
    IF (MAC address of mesh destination of data frame ==
Destination Mesh STA Address in Congestion Notification
element)
    THEN
        do not forward data frame until the corresponding
        Congestion Notification Expiration Timer of the
        corresponding access category has expired.
    ENDIF
ENDIF
```

[0060] The amended rule for postponing of data frame forwarding described in pseudo code above is generally described as follows:

[0061] The forwarding of a data frame at a receiving node is postponed for the case that:

the address of the next hop of the data frame is equal to the transmitter address of at least one received congestion control notification frame;

AND;

[0062] the destination address of the data frame is equal to at least one destination mesh node address of at least one

congestion notification element that has been received from the next hop of the data frame;

AND;

[0063] the corresponding Congestion Notification Expiration Timer for said next hop, destination address, and access category of the data frame has not expired.

[0064] An >>AND<< in the condition above is to be understood as logical AND.

[0065] In case that the destination field 204 of the at least one received congestion notification element is set to the broadcast address in order to apply the intra-mesh congestion control to all destinations that use the transmitter of the congestion notification element as next hop, the amended rule for postponing of data frame forwarding is generally described as follows:

[0066] The forwarding of a data frame at a receiving node is postponed for the case that:

the address of the next hop of the data frame is equal to the transmitter address of at least one received congestion control notification frame;

AND;

[0067] the destination mesh node address of at least one congestion notification element, that has been received from the next hop of the data frame, is equal to the broadcast address;

AND;

[0068] the corresponding Congestion Notification Expiration Timer for said next hop, destination address, and access category of the data frame has not expired.

[0069] The amended rule for forwarding of a data frame is generally described as follows:

[0070] A data frame is forwarded by a receiving node for the case that none of the two amended rules above for postponing the forwarding apply to the data frame.

[0071] The amended rule for forwarding the data frame is generally described as follows:

[0072] A data frame is forwarded by a receiving node for the case that:

the next hop address of the data frame is not equal to at least one next hop address from at least one received congestion control notification frame;

OR;

[0073] (the destination address of the data frame is not equal to at least one destination mesh node address from at least one congestion notification element that has been received from the next hop of the data frame; AND; each destination mesh node address from all congestion notification elements that have been received from the next hop of the data frame are different from the broadcast address);

OR;

[0074] all corresponding Congestion Notification Expiration Timers for said next hop, destination address, and access category of the data frame have expired.

[0075] These amended rules will not postpone the forwarding of data frames between mesh nodes which are not affected by a bottleneck, because the mesh destination is not to be congestion controlled. More specifically, according to FIG. 1,

this rule will not postpone traffic between the first mesh node A and the fifth mesh node E because the mesh destination, which is the fifth mesh node E, is not to be congestion controlled. The mesh destination address of the data frames is different from the address of the mesh destination signaled in the Congestion Notification element, i.e. the second, inner IF-condition is not fulfilled.

[0076] In the topology of FIG. 1, there is only a single dataflow DF1 initially affected by the bottleneck BN. FIG. 3 shows a topology where more than one data flow is affected by a bottleneck BN. Further on, affected data flows in FIG. 3 have different destinations.

[0077] According to FIG. 3, a plurality of mesh nodes A, B, C, D, E, F, G are logically interconnected by links which are depicted by lines interconnecting some of the mesh nodes A, B, C, D, E, F, G. A third dataflow DF3 is ranging from the first mesh node A passing the second mesh node B, the third mesh node C, the fourth mesh node D to a sixth mesh node F, which is the destination node F of said third dataflow DF3. A fourth dataflow DF4, illustrated by a dashed-and-dotted line, is ranging from the first mesh node A passing the second mesh node B, the third mesh node C, the fourth mesh node D to a seventh mesh node G, which is the destination node G of said fourth dataflow DF4. A fifth dataflow DF5, illustrated by a dotted line, is ranging from the first mesh node A passing the second mesh node B to the fifth mesh node E, which is the destination node E of said fifth dataflow DF5.

[0078] In a situation shown in FIG. 3, whereby the link between the third mesh node C and the fourth mesh node D becomes congested by a bottleneck BN, both, the third and the fourth dataflow DF3, DF4 contribute to the congestion at the bottleneck BN.

[0079] Hence, the third mesh node C has to send two Congestion Notification elements, one for each affected destination, when using the format of the Congestion Notification element as shown in FIG. 2.

[0080] According to an example embodiment it is proposed to include these two, or in general, at least one, Congestion Notification elements in a single Congestion Control Notification frame, the type of management frame of IEEE 802.11 that is used to transmit Congestion Notification elements, see Draft Standard IEEE P802.11s entitled: >>Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications—Mesh Networking<<, Version D7.03, November 2010, clause 7.4.15.6, entitled >>Congestion Control Notification frame format<<.

[0081] Alternatively, one could send two or more Congestion Control Notification frames with only one Congestion Notification element, but this will introduce an additional delay between the receptions of the different Congestion Notification elements due to the contention for the media access between the Congestion Control Notification frames.

[0082] FIG. 4 shows an exemplary structure of a congestion notification element with multiple destination addresses according to an alternative embodiment. The congestion notification element according to this embodiment uses a list of destinations in a single Congestion Notification element.

[0083] In analogy to the congestion notification element known from FIG. 2, the congestion notification element according to FIG. 4 includes a first field 400 captioned >>Element ID<< and having a length of one octet and a second field 402 captioned >>Length<< and having a length of one octet.

[0084] A third field 404 captioned >>Number of Destinations N<< defines a number (N) of fields captioned >>Desti-

nation Mesh STA Address \leqslant included within the Congestion Notification element. The third field **404** has a length of one octet.

[0085] Said fields **406**, **408**, **410** captioned \gg Destination Mesh STA Address \leqslant are adhered next to the third field **404**, namely a first destination address field **406** captioned \gg Destination Mesh STA Address #1 \leqslant , a field **408** in the drawing captioned $\gg \dots \leqslant$ and denoting a wildcard for a number N-2 of further destination address fields and, finally, an N-th destination address field **410** captioned \gg Destination Mesh STA Address #N \leqslant .

[0086] In analogy to the congestion notification element known from FIG. 2 the congestion notification element further contains four Congestion Notification Expiration Timer fields **412**, **414**, **416**, **418**. The Congestion Notification Expiration Timers are the same for all listed mesh destinations. If different Congestion Notification Expiration Timers are required for some mesh destinations, multiple Congestion Notification elements have to be used. A value of one in the third field **404** for the Number of Destinations field leads to an identical information of the Congestion Notification element shown in FIG. 2.

[0087] FIG. 5 shows an exemplary structure of a further amended congestion notification element according to a further embodiment. By using said notification element the usage of different Congestion Notification Expiration Timers per affected mesh destination is supported.

[0088] In analogy to the congestion notification element known from FIG. 2, the congestion notification element according to FIG. 5 includes a first field **500** captioned \gg Element ID \leqslant and having a length of one octet and a second field **502** captioned \gg Length \leqslant and having a length of one octet.

[0089] A third field **504** captioned \gg Number of Congestion Notifications N \leqslant defines a number (N) of Congestion Notifications **506**, **508**, **510** contained in this Congestion Notification element. The third field **504** has a length of one octet. [0090] Said Congestion Notifications **506**, **508**, **510** are adhered next to the third field **504**, namely a first Congestion Notification **506** captioned \gg Congestion Notification #1 \leqslant , a field **508** in the drawing captioned $\gg \dots \leqslant$ and denoting a wildcard for a number N-2 of further Congestion Notifications **508** and, finally, an N-th Congestion Notification **510** captioned \gg Congestion Notification #N \leqslant .

[0091] Each Congestion Notification **506**, **508**, **510** is defined by a number of sub-fields. Said sub-fields **550**, **552**, **554**, **556**, **558** of the first Congestion Notification **506** are shown in the drawing below the illustration of the congestion notification element.

[0092] A first sub-field **550** is captioned \gg Destination Mesh STA Address #1 \leqslant and has a length of six octets.

[0093] In analogy to the congestion notification element known from FIG. 2 the Congestion Notification further contains four Congestion Notification Expiration Timer sub-fields **552**, **554**, **556**, **558**.

[0094] Hence, a Congestion Notification consists of the address of the mesh destination for which congestion control has to be applied and the four Congestion Notification Expiration Timers. This allows different Congestion Notification Expiration Timers per affected mesh destination. A value of 1 for the Number of Congestion Notifications field **504** leads to an identical information of the Congestion Notification element shown in FIG. 2.

[0095] According to one embodiment, a representation of multiple congestion notifications is to transmit multiple Con-

gestion Notification elements according to or similar to FIG. 2 in a single Congestion Control Notification frame. This is the simplest representation of congestion notifications due to the small size of the Congestion Notification element and its easy processing.

[0096] In the following, an application example is described. A wireless mesh network including some data flows as shown in FIG. 3 is considered. There are three data flows:

[0097] data flow DF3 from mesh node A to mesh node F;

[0098] data flow DF4 from mesh node A to mesh node G;
and;

[0099] data flow DF5 from mesh node A to mesh node E.

[0100] The link from mesh node C to mesh node D is the bottleneck BN on the paths for data flows DF3, DF4.

[0101] Mesh node C becomes congested since more data frames for destination mesh nodes F,G arrive at mesh node C than mesh node C can actually forward to the next hop, mesh node D.

[0102] Data flows DF3, DF4 as the cause of the congestion can be derived by inspecting the respective data frames, wherein the transmitter address is set to mesh node B and the mesh destination address is set to mesh node F or mesh node G, respectively and from the forwarding information of the paths, i.e. precursor mesh node for destination mesh nodes F, G.

[0103] Mesh node C constructs two Congestion Notification elements according to an embodiment as shown in FIG. 6 and FIG. 7. The structure of said Congestion Notification elements in FIG. 6 and FIG. 7 corresponds to the structure according to FIG. 2.

[0104] The congestion notification element according to FIG. 6 includes a first field **600** captioned \gg Element ID \leqslant . It contains a value that indicates a Congestion Notification element. A second field **602** is captioned \gg Length \leqslant . The second field **602** carries a decimal value of 14. A third field **604** specifying the destination mesh node address carries the MAC address of the mesh node F.

[0105] The congestion notification element contains four Congestion Notification Expiration Timer fields **606**, **608**, **610**, **612**, namely a first Congestion Notification Expiration Timer field **606** captioned \gg Congestion Notification Expiration Timer AC_BK \leqslant , a second Congestion Notification Expiration Timer field **608** captioned \gg Congestion Notification Expiration Timer AC_BE, a third Congestion Notification Expiration Timer field **610** captioned \gg Congestion Notification Expiration Timer AC_VI \leqslant and a fourth Congestion Notification Expiration Timer field **612** captioned \gg Congestion Notification Expiration Timer AC_VO \leqslant .

[0106] The first Congestion Notification Expiration Timer field **606** carries an exemplary value of t0, the second Congestion Notification Expiration Timer field **608** carries an exemplary value of t1, the third Congestion Notification Expiration Timer field **610** carries an exemplary value of t2 and the fourth Congestion Notification Expiration Timer field **612** carries an exemplary value of t3. However, the actual values of the four Congestion Notification Expiration Timers are not necessary for the explanation hereinafter.

[0107] The congestion notification element according to FIG. 7 includes a first field **700** captioned \gg Element ID \leqslant . It contains a value that indicates a Congestion Notification element. A second field **702** is captioned \gg Length \leqslant . The second field **702** carries a decimal value of 14. A third field **704**

specifying the destination mesh node address carries the MAC address of the mesh node G.

[0108] The congestion notification element contains four Congestion Notification Expiration Timer fields **706**, **708**, **710**, **712**, namely a first Congestion Notification Expiration Timer field **706** captioned >>Congestion Notification Expiration Timer AC_BK≤≤, a second Congestion Notification Expiration Timer field **708** captioned >>Congestion Notification Expiration Timer AC_BE, a third Congestion Notification Expiration Timer field **710** captioned >>Congestion Notification Expiration Timer AC_VI≤≤ and a fourth Congestion Notification Expiration Timer field **712** captioned >>Congestion Notification Expiration Timer AC_VO≤≤.

[0109] The first Congestion Notification Expiration Timer field **706** carries an exemplary value of t4, the second Congestion Notification Expiration Timer field **708** carries an exemplary value of t5, the third Congestion Notification Expiration Timer field **710** carries an exemplary value of t6 and the fourth Congestion Notification Expiration Timer field **712** carries an exemplary value of t7. However, the actual values of the four Congestion Notification Expiration Timers are not necessary for the explanation hereinafter.

[0110] The Congestion Notification elements according to FIGS. 6 and 7 are transmitted in a single Congestion Control Notification frame to the precursor on the paths of data flows DF3,DF4, which is mesh node B.

[0111] After mesh node B has received the Congestion Control Notification frame with the two Congestion Notification elements according to FIG. 6 and FIG. 7, said mesh node B stores the relevant information of each congestion notification, including:

[0112] the mesh node, or next hop, to which mesh node B should postpone the forwarding of data frames. This is the transmitter of the Congestion Control Notification frame. Its address is contained in the transmitter address field of the header of the Congestion Control Notification frame.

[0113] the mesh destination for which mesh node B should postpone the forwarding of data frames if the next hop corresponds to the sender of the congestion notification. This address is contained in the Destination Mesh STA Address field of the Congestion Notification element.

[0114] the four congestion notification expiration timers, for each access category.

[0115] In the example, the stored values of Congestion Notifications at mesh node B would be stored according to a Table 1 as shown below:

next hop	mesh destination	congestion notification expiration timer AC_BK	congestion notification expiration timer AC_BE	congestion notification expiration timer AC_VI	congestion notification expiration timer AC_VO
mesh node	mesh node F	t0	t1	t2	t3
node C	mesh node G	t4	t5	t6	t7

A » mesh STA « (» mesh station «) is a frequently used expression for a mesh node as used herein.

[0116] Mesh node B now receives a data frame of data flow DF3, that is, a data frame with mesh node F as mesh destination. Mesh node B checks the conditions for postponing the forwarding of the data frame:

[0117] The first IF-condition is true, because the MAC address of the next hop of the data frame (according to the forwarding information for mesh node F, this is mesh node C) corresponds to the MAC address of the transmitter of a Congestion Notification element. Note that Table 1 shown above includes two entries for mesh node C.

[0118] The second IF-condition is true, because the mesh destination MAC address of the data frame, which is mesh node F, corresponds to a mesh destination entry of Table 1 shown above for mesh node C as next hop.

[0119] Since both nested IF-conditions are true, the forwarding of the data frame for mesh destination F is postponed.

[0120] Mesh node B now receives a data frame of data flow DF4, that is, a data frame with mesh node G as mesh destination. Mesh node B checks the conditions for postponing the forwarding of the data frame:

[0121] The first IF-condition is true, because the MAC address of the next hop of the data frame (according to the forwarding information for mesh node G, this is mesh node C) corresponds to the MAC address of the transmitter of a Congestion Notification element. Note that Table 1 shown above includes two entries for mesh node C.

[0122] The second IF-condition is true, because the mesh destination MAC address of the data frame, which is mesh node G, corresponds to a mesh destination entry of Table 1 shown above for mesh node C as next hop.

[0123] Since both nested IF-conditions are true, the forwarding of the data frame for mesh destination G is postponed.

[0124] Mesh node B now receives a data frame of data flow DF5, that is, a data frame with mesh node E as mesh destination. Mesh node B checks the conditions for postponing the forwarding of the data frame:

[0125] The first IF-condition is false, because there is no entry for the MAC address of the next hop of the data frame (according to the forwarding information for mesh node G, this is mesh node E) in column "next hop" of the stored congestion notifications.

[0126] Since already the first IF-condition is false, the forwarding of the data frame for mesh destination E is not postponed. The data frame is, therefore, forwarded.

[0127] Assume, that mesh node B becomes congested as well, because it does not forward data frames for destinations

F and G that it receives. In the same way as mesh node C, it will construct two Congestion Notification elements which are sent by a single Congestion Control Notification frame to mesh node A.

[0128] The content of the Destination Mesh node Address fields is the same, only the times for the Congestion Notification Expiration Timer fields might be different. Table 2 below shows the stored values for the received congestion notifications at mesh node A.

[0129] In the example, the stored values of Congestion Notifications at mesh node A would be stored according to Table 2 as shown below:

next hop	mesh destination	congestion notification expiration timer AC_BK	congestion notification expiration timer AC_BE	congestion notification expiration timer AC_VI	congestion notification expiration timer AC_VO
mesh node	mesh node F
node B	mesh node G

[0130] Mesh node A now wants to forward a data frame of data flow DF3, that is, a data frame with mesh node F as mesh destination. Mesh node A checks the conditions for postponing the forwarding of the data frame:

[0131] The first IF-condition is true, because the MAC address of the next hop of the data frame (according to the forwarding information for mesh node F, this is mesh node B) corresponds to the MAC address of the transmitter of a Congestion Notification element. Note that Table 2 shown above includes two entries for mesh node B.

[0132] The second IF-condition is true, because the mesh destination MAC address of the data frame (which is mesh node F) corresponds to a mesh destination entry of Table 2 for mesh node B as next hop.

[0133] Since both nested IF-conditions are true, the forwarding of the data frame for the destination mesh node F is postponed.

[0134] Mesh node A now wants to forward a data frame of data flow DF4, that is, a data frame with mesh node G as mesh destination. Mesh node A checks the conditions for postponing the forwarding of the data frame:

[0135] The first IF-condition is true, because the MAC address of the next hop of the data frame (according to the forwarding information for mesh node G, this is mesh node B) corresponds to the MAC address of the transmitter of a Congestion Notification element. Note that Table 2 shown above includes two entries for mesh node B.

[0136] The second IF-condition is true, because the mesh destination MAC address of the data frame (which is mesh node G) corresponds to a mesh destination entry of Table 2 for mesh node B as next hop.

[0137] Since both nested IF-conditions are true, the forwarding of the data frame for destination mesh node G is postponed.

[0138] Mesh node A now wants to forward a data frame of data flow DF5, that is, a data frame with mesh node E as mesh

destination. Mesh node A checks the conditions for postponing the forwarding of the data frame:

[0139] The first IF-condition is true, because the MAC address of the next hop of the data frame (according to the forwarding information for mesh node E, this is mesh node B) corresponds to the MAC address of the transmitter of a Congestion Notification element. Note that Table 2 shown above includes two entries for mesh node B.

[0140] The second IF-condition is false, because there is no entry for the mesh destination MAC address of the

data frame (which is mesh node E) in column >>mesh destination<= of Table 2 with mesh node B as next hop.

[0141] Since not both of the nested IF-conditions are true, the forwarding of the data frame for mesh destination E is not postponed. The data frame is forwarded. This correct and advantageous behavior at mesh node A is possible with the disclosed embodiments.

What is claimed is:

1. A congestion notification element for indicating a congestion status of a mesh node in a mesh network, the congestion notification element comprising:

at least one field specifying a mesh destination for applying an intra-mesh congestion control.

2. Congestion notification element according to claim 1, wherein the mesh destination is determined by an address of a destination mesh node.

3. Congestion notification element according to claim 1, wherein the mesh destination is determined by a broadcast address.

4. Congestion notification element according to claim 1, further comprising field specifying a number of fields, each specifying a respective mesh destination.

5. Congestion notification element according to claim 1, comprising at least one congestion notification, each congestion notification including

one of the at least one field specifying the mesh destination, at least one field specifying a congestion notification expiration timer, and

a field specifying a number of congestion notifications included in the congestion notification element.

6. A congestion control notification frame, Comprising: at least one congestion notification element, each congestion notification element comprising:

at least one field specifying a mesh destination for applying an intra-mesh congestion control.

7. A congestion control notification frame according to claim 6, comprising fields defined by IEEE P802.11s.

8. A method of controlling congestion in a mesh network, the method comprising:

receiving at a receiving mesh node from a transmitting mesh node at least one congestion control notification frame including at least one congestion notification element;

reading at least one transmitter mesh node address of said at least one congestion control notification frame as a next hop address;

reading at least one destination mesh node address of said at least one congestion notification element, said destination mesh node address being an individual address or a broadcast address; and

reading at least one congestion notification expiration timer of said at least one congestion notification element.

9. The method according to claim **8**, comprising:
receiving at said receiving mesh node a data frame including a destination address of said data frame;
reading the destination address of said data frame; and
reading a next hop address for the destination address of said data frame from a forwarding information.

10. The method according to claim **9**, comprising:
postponing a forwarding of said data frame by said receiving node for a case in which said next hop address for said data frame is equal to at least one next hop address from said at least one congestion control notification frame;
wherein said destination address of said data frame is equal to at least one destination mesh node address from said at least one next hop address of said at least one congestion notification element.

11. The method according to claim **9**, comprising:
postponing a forwarding of said data frame by said receiving node for a case in which said next hop address for said data frame is equal to at least one next hop address from said at least one congestion control notification frame;
wherein at least one destination mesh node address from said at least one next hop address of said at least one congestion notification element is equal to said broadcast address.

12. The method according to claim **10**, wherein said postponing of said forwarding of said data frame is performed until a corresponding congestion notification expiration timer of a corresponding access category has expired.

13. A mesh node in a mesh network comprising non-transitory computer-readable media storing instructions executable by a processor to:

receive from a transmitting mesh node at least one congestion control notification frame including at least one congestion notification element;

read at least one transmitter mesh node address of said at least one congestion control notification frame as a next hop address;

read at least one destination mesh node address of said at least one congestion notification element, said destination mesh node address being an individual address or a broadcast address; and

read at least one congestion notification expiration timer of said at least one congestion notification element.

14-17. (canceled)

18. The mesh node according to claim **13**, wherein the instructions are further executable to:
receive at said mesh node a data frame including a destination address of said data frame;
read the destination address of said data frame; and
read a next hop address for the destination address of said data frame from a forwarding information.

19. The mesh node according to claim **18**, wherein the instructions are further executable to:
postpone a forwarding of said data frame by said receiving node for a case in which said next hop address for said data frame is equal to at least one next hop address from said at least one congestion control notification frame;
wherein said destination address of said data frame is equal to at least one destination mesh node address from said at least one next hop address of said at least one congestion notification element.

20. The method according to claim **18**, wherein the instructions are further executable to:
postpone a forwarding of said data frame by said receiving node for a case in which said next hop address for said data frame is equal to at least one next hop address from said at least one congestion control notification frame;
wherein at least one destination mesh node address from said at least one next hop address of said at least one congestion notification element is equal to said broadcast address.

21. The mesh node according to claim **19**, wherein said postponing of said forwarding of said data frame is performed until a corresponding congestion notification expiration timer of a corresponding access category has expired.

* * * * *