(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
02 February 2023 (02.02.2023)

WIPO | PCT

(10) International Publication Number
WO 2023/006167 A1

(51) International Patent Classification:
H04L 45/24 (2022.01) *H04L 45/64* (2022.01)
H04L 45/125 (2022.01)

(21) International Application Number:
PCT/EP2021/070769

(22) International Filing Date:
26 July 2021 (26.07.2021)

(25) Filing Language: English

(26) Publication Language: English

(71) Applicant: **HUAWEI TECHNOLOGIES CO., LTD.** [CN/CN]; Huawei Administration Building Bantian Longgang District, Shenzhen, Guangdong 518129 (CN).

(72) Inventor; and
(71) Applicant *(for US only)*: **LIASKOS, Christos** [GR/DE]; c/o Huawei Technologies Duesseldorf GmbH, Riesstr. 25, 80992 Munich (DE).

(72) Inventors: **ALEXANDRIS, Konstantinos**; c/o Huawei Technologies Duesseldorf GmbH, Riesstr. 25, 80992 Munich (DE). **TROSSEN, Dirk**; c/o Huawei Technologies Duesseldorf GmbH, Riesstr. 25, 80992 Munich (DE). **TANG, Siyu**; c/o Huawei Technologies Duesseldorf GmbH, Riesstr. 25, 80992 Munich (DE).

(74) Agent: **KREUZ, Georg**; Huawei Technologies Duesseldorf GmbH, Riesstr. 25, 80992 Munich (DE).

(81) Designated States *(unless otherwise indicated, for every kind of national protection available)*: AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, IT, JO, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, WS, ZA, ZM, ZW.

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

Published:
— *with international search report (Art. 21(3))*

(54) **Title:** NETWORK TRAFFIC ENGINEERING BASED ON TRAVERSING DATA RATE

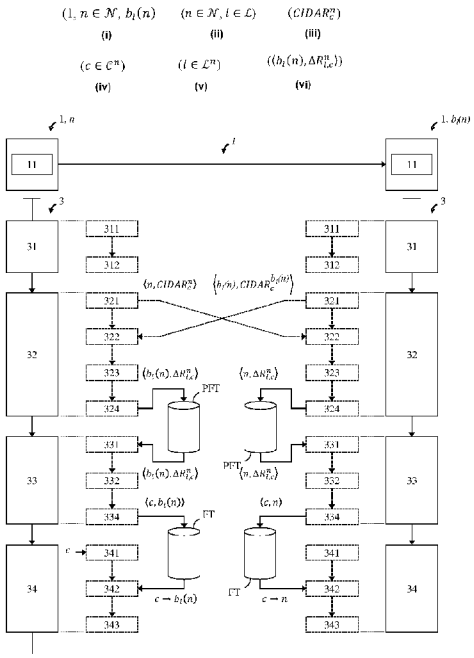


FIG. 2

(57) **Abstract:** Disclosed is a network node (i) of a network (ii). The network node (1) comprises a processor (11) configured to: determine (31) respective continuously integrated data rate (iii) metrics for a plurality of communication relations (iv) between the network node (1, n) and one or more neighbor network nodes (1, $b_l(n)$) connected to the network node (1, n) via respective network links (v); establish (32) respective potential forwarding rules (vi) for the plurality of communication relations (c) based on the respective determined continuously integrated data rate (iii) metrics; allocate (33) the plurality of communication relations (c) to the respective network links (l) in accordance with those of the respective potential forwarding rules (vi) which increase a bandwidth share of the corresponding associated relation (c) on the respective network link (l); and perform (34) a forwarding for the plurality of communication relations (c) in accordance with the allocation. This provides a throughput-optimizing traffic engineering technique which is unaffected by indications of congestion such as build-up of queues.

WO 2023/006167 A1

NETWORK TRAFFIC ENGINEERING BASED ON TRAVERSING DATA RATE

TECHNICAL FIELD

The present disclosure generally relates to end-to-end traffic engineering in packet-switched networks. In particular, the present disclosure specifically relates to a network node, a controller node, and associated methods of operating the same.

BACKGROUND ART

In a typical scenario, a service client C may issue a request for service SF, to be served by any one of multiple service instances SF_1, …, SF_i deployed at various network locations, and thereby constituting a communication relation c.

A stateful service SF may require that once a communication relation c has been established to a particular service instance SF_i, subsequent communications pertaining to the communication relation c must consistently be served by this particular service instance SF_i due to state being built up (a.k.a. instance affinity). By contrast, a stateless service SF such as file retrieval for over-the-top (OTT) video may be provided by ever-changing service instances SF_i in response to each request.

In any case, mapping service requests involves issues such as how to select the 'best' service instance to serve the client's request, and how to distribute requests from more than one client (or the resulting communication relations c) in a manner that compute/network capabilities are best used.

SUMMARY

At least some of the foregoing and other issues are solved by the features of the independent claims. Further implementation forms are apparent from the dependent claims, the description and the figures.

According to a first aspect, a network node of a network is provided. The network node comprises a processor configured to: determine respective continuously integrated data rate (*CIDAR*) metrics for a plurality of communication relations between the network node and one or more neighbour network nodes connected to the network node via respective network links;

establish respective potential forwarding rules for the plurality of communication relations based on the respective determined continuously integrated data rate metrics; allocate the plurality of communication relations to the respective network links in accordance with those of the respective potential forwarding rules which increase a bandwidth share of the corresponding associated relation on the respective network link; and perform a forwarding for the plurality of communication relations in accordance with the allocation.

Allocating the plurality of communication relations to network links in view of increased bandwidth shares between themselves and their neighbours of relevance effectively means that respective network nodes are enabled to perform traffic engineering (TE). TE may herein be used to refer to techniques for mapping the given plurality of communication relations to the given network topology with the objective to enhance one or more performance metrics of the network, such as network throughput, for example. In response to a reporting of CIDAR metrics from its neighbour nodes, the respective network node determines its 'best' next hops and updates its forwarding table. As such, the forwarding decision is not done at the data plane but as a control plane operation, and forwarding is reduced to using the latest update to the forwarding table, which is compatible with typical 'next hop' forwarding actions.

Determining respective *CIDAR* metrics for the plurality of communication relations between the neighbouring network nodes decouples the subsequent forwarding from any indication of congestion such as a build-up of queues. The latter would counteract an objective of reliable transmission protocols such as the transmission control protocol (TCP) to reduce queue build-up in response to congestion.

Establishing respective potential forwarding rules for the plurality of communication relations based on the determined respective CIDAR metrics couples possible TE (and thus subsequent forwarding) decisions to corresponding CIDAR metrics, which quantify the consequences of choosing the possible TE / forwarding rule in connection with the objective to enhance the one or more performance metrics of the network.

Performing the forwarding for the plurality of communication relations in accordance with the allocation enforces instance affinity, i.e., scheduling applied at transaction, not packet level.

The neighbour-centric TE technique is topology-agnostic, as it may be applied directly in ingress-egress topologies, such as in leaf-spine topologies in data centres (DC), or applied recursively between ingress and egress nodes of a fully routed architecture (e.g., service routing network), by treating the egress as ingress towards its next hop unless the next hop is the egress towards one or more service instances. As such, the TE technique may be applied on a network-wide / end-to-end (E2E) basis.

The TE technique may be performed by respective (ingress and possibly intermediate) network nodes and thus mandates no centralized entity, but may still revert to a centralized TE server, if needed (see controller node below).

The TE technique ensures compatibility with any packet-level scheduler, as establishing and applying forwarding rules involves the network layer only and does not affect packet dispatching further.

In a possible implementation form, the processor may further be configured to, so as to determine the respective continuously integrated data rate metrics for the plurality of communication relations between the network node and the one or more neighbour network nodes connected to the network node via the respective network links: measure a respective pass-through data rate of the plurality of communication relations at periodic time instants; and integrate the respective measured data rate over a continuous time period.

CIDAR metrics take into account a pass-through / traversing data rate only, and not a build-up of queues, so that using CIDAR metrics optimizes network performance also in connection with TCP traffic.

In a possible implementation form, the processor may further be configured to, so as to establish the respective potential forwarding rules for the plurality of communication relations based on the respective determined continuously integrated data rate metrics: advertise the determined respective continuously integrated data rate metrics in association with an identifier of the network node; receive respective continuously integrated data rate metrics from the one or more neighbour network nodes respectively in association with an identifier of the respective neighbour network node; calculate differential metrics of the determined respective continuously integrated data rate metrics of the network node and the received respective

continuously integrated data rate metrics of the respective neighbour network node; and populate a preference table of the network node with the respective potential forwarding rules for the plurality of communication relations, the respective potential forwarding rule comprising an identifier of a respective neighbour network node and the calculated differential metric associated with the respective neighbour network node.

The TE technique uses differences of CIDAR metrics of two neighbouring nodes during some time interval, here termed calculated differential metrics, as potential forwarding metrics. This may outperform other techniques such as backpressure-based techniques, for example, which require a build-up of packet queues. As such, congestion may be avoided proactively, making the TE technique an appropriate solution for DCs characterized by minimal latency and/or variable bit rate (VBR) traffic.

In a possible implementation form, the processor may further be configured to, so as to advertise the determined respective continuously integrated data rate metrics: advertise the determined respective continuously integrated data rate metrics via an interior gateway protocol or an exterior gateway protocol.

The TE technique may thus re-use and integrate with existing protocol infrastructure. A service instance (with publicly routable IP address) acting as egress may report / advertise its CIDAR metrics directly, while a client acting as an ingress may receive reports / advertisements in addition to other network nodes. Both variants add architectural flexibility. The TE technique may further differentiate between the plurality of communication relations, where only specific communication relations effectively participate. This provides architectural scalability.

In a possible implementation form, the processor may further be configured to, so as to establish the respective potential forwarding rules for the plurality of communication relations based on the respective determined continuously integrated data rate metrics: send the determined respective continuously integrated data rate metrics to a controller node of the network.

The TE technique may revert to a centralized TE server (i.e., controller node), if needed. If so, the server may establish potential forwarding rules for the ingress nodes (or all nodes in a fully routed architecture). This may effectively off-load the network nodes.

In a possible implementation form, the processor may further be configured to, so as to send the determined respective continuously integrated data rate metrics to the controller node of the network: send the determined respective continuously integrated data rate metrics to the controller node in response to a polling indication received from the same.

The TE technique may revert to a centralized TE server, if needed. If so, the server may poll the CIDAR metrics alternatively to periodic and/or occasional reporting by the network nodes. This may effectively limit a signalling load of the network and the controller node.

In a possible implementation form, the processor may further be configured to, so as to allocate the plurality of communication relations to the respective network links in accordance with those of the respective potential forwarding rules which increase a bandwidth share of the corresponding associated relation on the respective network link: sort the potential forwarding rules of the network node by descending value of the calculated differential metric; mark, in accordance with the descending order of the sorted potential forwarding rules, one of the potential forwarding rules for the associated communication relation for deployment if it improves a throughput of the network; and populate a forwarding table of the network node with the respective entries for the potential forwarding rule marked for deployment.

The TE technique may use this throughput-optimizing heuristic that favours high values of the calculated differential metric to obtain an acceptable TE solution within an acceptable timeframe.

In a possible implementation form, the processor may further be configured to, so as to allocate the plurality of communication relations to the respective network links in accordance with those of the respective potential forwarding rules which increase a bandwidth share of the corresponding associated relation on the respective network link: receive respective entries for the forwarding table of the network node for those of the potential forwarding rules marked for deployment, the respective forwarding entry comprising an identifier of the communication relation associated with the marked potential forwarding rule and the neighbour network node of the marked potential forwarding rule; and populate the forwarding table of the network node with the respective entries for those of the potential forwarding rules marked for deployment.

The TE technique may revert to a centralized TE server (i.e., controller node), if needed. If so, the respective network node may receive ready-made entries for population of the local forwarding table. This may effectively off-load the network nodes.

In a possible implementation form, the processor may further be configured to, so as to perform the forwarding for the plurality of communication relations in accordance with the allocation: receive a packet associated with a particular one of the one or more communication relations; retrieve the forwarding entry associated with the particular communication relation from the forwarding table; and forward the received packet in accordance with the retrieved forwarding entry.

According to a second aspect, a controller node of a network is provided, comprising a processor configured to: establish respective potential forwarding rules for a plurality of communication relations between one or more network nodes of the network and respective one or more neighbour network nodes connected to the network nodes via respective network links based on respective determined continuously integrated data rate metrics; and allocate the plurality of communication relations to the respective network links in accordance with those of the respective potential forwarding rules which increase a bandwidth share of the associated communication relation on the respective network link.

The TE technique may revert to a centralized TE server (i.e., controller node), if needed. If so, the server may effectively off-load the network nodes with respect to the establishment of respective potential forwarding rules, and also with respect to the mapping (i.e., TE) of the plurality of communication relations to respective network links.

In a possible implementation form, the processor may further be configured to, so as to establish the respective potential forwarding rules for the plurality of communication relations between the one or more network nodes of the network and respective one or more neighbour network nodes connected to the network nodes via respective network links based on respective determined continuously integrated data rate metrics: receive respective continuously integrated data rate metrics from the one or more network nodes respectively in association with an identifier of the respective neighbour network node; calculate differential metrics of the received respective continuously integrated data rate metrics of the one or more network nodes and the received respective continuously integrated data rate metrics of the respective one or

more neighbour network nodes; and populate respective preference tables for the one or more network nodes with the respective potential forwarding rules for the respective plurality of communication relations, the respective potential forwarding rule comprising an identifier of a respective one of the one or more neighbour network nodes and the calculated differential metric associated with the respective one of the one or more neighbour network nodes.

The TE technique may revert to a centralized TE server (i.e., controller node), if needed. If so, the server may use differences of CIDAR metrics of two neighbouring nodes during some time interval, here termed calculated differential metrics, as potential forwarding metrics. This may outperform other techniques such as backpressure-based techniques, for example, which require a build-up of packet queues. As such, congestion may be avoided proactively, making the TE technique an appropriate solution for DCs characterized by minimal latency and/or variable bit rate (VBR) traffic.

In a possible implementation form, the processor may further be configured to, so as to receive the respective continuously integrated data rate metrics from the one or more network nodes: receive respective continuously integrated data rate metrics from the one or more network nodes in response to a polling indication sent to the same.

The TE technique may revert to a centralized TE server (i.e., controller node), if needed. If so, the server may poll the CIDAR metrics alternatively to periodic and/or occasional reporting by the network nodes. This may effectively limit a signalling load of the network and the controller node.

In a possible implementation form, the processor may further be configured to, so as to allocate the plurality of communication relations to the respective network links in accordance with those of the respective potential forwarding rules which increase a bandwidth share of the associated communication relation on the respective network link: sort the potential forwarding rules of the respective network node by descending value of the calculated differential metric; mark, in accordance with the descending order of the sorted potential forwarding rules, one of the potential forwarding rules for the associated communication relation for deployment if it improves a throughput of the network; and send respective entries for a forwarding table of the respective network node for the potential forwarding rule marked for deployment, the respective entry comprising an identifier of the communication relation associated with the

marked potential forwarding rule and the neighbour network node of the marked potential forwarding rule.

The TE technique may revert to a centralized TE server (i.e., controller node), if needed. If so, the server may use the same throughput-optimizing heuristic that favours high values of the calculated differential metric to obtain an acceptable TE solution within an acceptable timeframe.

In a possible implementation form, the plurality of communication relations may respectively comprise one of: a destination-based communication relation, a flow-based communication relation, a path-based communication relation, a service-based communication relation, and a label-based communication relation.

The TE technique may thus be applied to a variety of communication relations.

According to a third aspect, a method of operating a network node of a network is provided. The method comprises: determining respective continuously integrated data rate metrics for a plurality of communication relations between the network node and one or more neighbour network nodes connected to the network node via respective network links; establishing respective potential forwarding rules for the plurality of communication relations based on the respective determined continuously integrated data rate metrics; allocating the plurality of communication relations to the respective network links in accordance with those of the respective potential forwarding rules which increase a bandwidth share of the corresponding associated relation on the respective network link; and performing a forwarding for the plurality of communication relations in accordance with the allocation.

In a possible implementation form, the method may be performed by the network node of the first aspect or any of its embodiments.

According to a fourth aspect, a method of operating a controller node of a network is provided. The method comprises: establishing respective potential forwarding rules for a plurality of communication relations between one or more network nodes of the network and respective one or more neighbour network nodes connected to the network nodes via respective network links based on respective determined continuously integrated data rate metrics; and allocating

the plurality of communication relations to the respective network links in accordance with those of the respective potential forwarding rules which increase a bandwidth share of the associated communication relation on the respective network link.

In a possible implementation form, the method may be performed by the controller node of the second aspect or any of its embodiments.

According to a fifth aspect, a computer program is provided, comprising executable instructions which, when executed by a processor, cause the processor to perform the method of any one of the third and fourth aspects or any embodiment thereof.

The technical effects and advantages described above in relation with the network node and the controller node equally apply to the corresponding methods of operating the same having corresponding features.

BRIEF DESCRIPTION OF DRAWINGS

The above-described aspects and implementations will now be explained with reference to the accompanying drawings, in which the same or similar reference numerals designate the same or similar elements.

The drawings are to be regarded as being schematic representations, and elements illustrated in the drawings are not necessarily shown to scale. Rather, the various elements are represented such that their function and general purpose become apparent to those skilled in the art.

FIG. 1 illustrates an exemplary network scenario in accordance with the present disclosure;

FIG. 2 illustrates two neighbouring network nodes and a corresponding method of operating the same in accordance with the present disclosure; and

FIG. 3 illustrates a network node, a controller node and corresponding methods of operating the same in accordance with the present disclosure.

DETAILED DESCRIPTIONS OF DRAWINGS

In the following description, reference is made to the accompanying drawings, which form part of the disclosure, and which show, by way of illustration, specific aspects of embodiments of the invention or specific aspects in which embodiments of the present invention may be used. It is understood that embodiments of the invention may be used in other aspects and comprise structural or logical changes not depicted in the figures. The following detailed description, therefore, is not to be taken in a limiting sense, and the scope of the present invention is defined by the appended claims.

For instance, it is understood that a disclosure in connection with a described method may also hold true for a corresponding apparatus or system configured to perform the method and vice versa. For example, if one or a plurality of specific method steps are described, a corresponding device may include one or a plurality of units, e.g. functional units, to perform the described one or plurality of method steps (e.g. one unit performing the one or plurality of steps, or a plurality of units each performing one or more of the plurality of steps), even if such one or more units are not explicitly described or illustrated in the figures. On the other hand, for example, if a specific apparatus is described based on one or a plurality of units, e.g. functional units, a corresponding method may include one step to perform the functionality of the one or plurality of units (e.g. one step performing the functionality of the one or plurality of units, or a plurality of steps each performing the functionality of one or more of the plurality of units), even if such one or plurality of steps are not explicitly described or illustrated in the figures. Further, it is understood that the features of the various exemplary embodiments and/or aspects described herein may be combined with each other, unless specifically noted otherwise.

FIG. 1 illustrates an exemplary network scenario in accordance with the present disclosure.

The scenario comprises a communication network indicated by a cloud shape and comprising network nodes 1 interconnected by network layer (i.e., layer-3) links.

Network nodes 1 as used herein may generally relate to network nodes suitable for layer-3 routing. In other words, the communication network may be a routed network, such as an IP network.

The network nodes 1 are required for distributed TE in accordance with the method 3 of operating a network node 1 which will be explained in more detail in connection with FIG. 2.

The scenario may further comprise a controller node 2 for centralized TE in accordance with the methods 3, 4 which will be explained in more detail in connection with FIG. 3.

Attached to the communication network are service clients C deployed at various network locations, and multiple instances SF_1, ..., SF_I of a service function SF deployed at servers S at various network locations as well.

In this exemplary scenario, any of the service clients C may issue a request for service SF, to be served by any one of the multiple service instances SF_i. Such a mapping involves issues such as (i) how to select the 'best' service instance to serve the client's request, and (ii) how to distribute service requests from more than one client in a manner that compute/network capabilities are best used.

In accordance with the present disclosure, a TE-based solution of especially the latter issue may be achieved, either among the network nodes 1 (distributed TE implementation, see FIG. 2), or between the network nodes 1 and a controller node 2 (centralized TE implementation, see FIG. 3), as will be explained in more detail below.

FIG. 2 illustrates two neighbouring network nodes 1 of a network ($n \in \mathcal{N}$, $l \in \mathcal{L}$) and a corresponding method 3 of operating these network nodes 1 in accordance with the present disclosure.

The distributed TE implementation may involve neighboring network nodes as follows:

The network node 1, $n \in \mathcal{N}$ depicted at the top left of FIG. 2 and the network node 1, $b_l(n) \in \mathcal{N}$ shown at the top right of FIG. 2 are connected via a network link l, which is illustrated as a directed network link without loss of generality. The network nodes 1, n, $b_l(n)$ thus have a neighborship relation in that the network node 1, $b_l(n)$ is a neighbour of the network node 1, n (and vice versa in case of an additional, oppositely directed network link). The expression $b_l(n)$ may thus be understood as meaning "neighbor of network node n via/across network link l".

This notion of neighborship relation maps directly to single-hop ingress-egress topologies, such as leaf-spine topologies, in which multiple leaf nodes $b_l(n)$ are connected to a same spine node n via respective links l.

Furthermore, the neighborship relation may also be applied recursively between ingress and egress nodes of fully routed multi-hop topologies, by treating each egress as ingress towards its next hop unless the next hop is the egress towards one or more of the multiple service instances SF_1, …, SF_i. That is to say, an end-to-end communication relation between a service client C and a service instance SF_i across the network may generally involve an ingress node 1, n connected to the service client C, an egress node 1, $b_l(n)$ connected to the service instance SF_i and a number of intermediate network nodes 1 respectively acting as an egress node $b_l(n)$ towards the service client C (*"upstream"*) and further acting as an ingress node n towards the service instance SF_i (*"downstream"*).

The network nodes 1, n, $b_l(n)$ of FIG. 2 comprise a respective processor 11 configured to perform a respective instance of a method 3 of operating a network node 1, which is depicted below the respective network node 1 in different levels of detail.

The method 3 comprises in its most concise representation shown at the side edges of FIG. 2: determining **31** respective continuously integrated data rate $CIDAR_c^n$ metrics for a plurality of communication relations $c \in \mathcal{C}^n$ between the network node 1, n and one or more neighbor network nodes 1, $b_l(n)$ connected to the network node 1, n via respective network links $l \in \mathcal{L}^n$; establishing **32** respective potential forwarding rules $\langle b_l(n), \Delta R_{l,c}^n \rangle$ for the plurality of communication relations c based on the respective determined continuously integrated data rate $CIDAR_c^n$ metrics; allocating **33** the plurality of communication relations c to the respective network links l in accordance with those of the respective potential forwarding rules $\langle b_l(n), \Delta R_{l,c}^n \rangle$ which increase a bandwidth share of the corresponding associated relation c on the respective network link l; and performing **34** a forwarding for the plurality of communication relations c in accordance with the allocation.

As will be appreciated by those skilled in the art, the respective network node 1, n, $b_l(n)$ comprises, in its most concise representation, device features corresponding to the above-

identified method steps 31 – 34, and a device feature of the network node 1, n, $b_l(n)$ and its corresponding method feature/step of the method 3 represents a same functionality.

In a more extensive presentation of the respective network node 1, n, $b_l(n)$, the processor 11 may further be configured to, so as to determine **31** the respective continuously integrated data rate $CIDAR_c^n$ metrics for the plurality of communication relations $c \in \mathcal{C}^n$ between the network node 1, n and the one or more neighbor network nodes 1, $b_l(n)$ connected to the network node 1, n via the respective network links $l \in \mathcal{L}^n$: measure **311** a respective pass-through data rate $U_{n,c}^{(k-1)}$ of the plurality of communication relations c at periodic time instants $k \cdot \Delta t$ (where $k \in [1; K]$); and integrate **312** the respective measured data rate $CIDAR_c^n = \sum_{k=1}^{K} U_{n,c}^{(k-1)}$ over a continuous time period $K \cdot \Delta t$.

Contrary to what the illustration in FIG. 2 may suggest, this measurement and integration is performed periodically to ensure that the continuously integrated data rate $CIDAR_c^n$ metrics are always accurate.

The processor 11 may further be configured to, so as to establish **32** the respective potential forwarding rules $\langle b_l(n), \Delta R_{l,c}^n \rangle$ for the plurality of communication relations c based on the respective determined continuously integrated data rate $CIDAR_c^n$ metrics: advertise **321** the determined respective continuously integrated data rate $CIDAR_c^n$ metrics in association with an identifier of the network node 1, n (of note, each network node 1 considers itself as node 1, n when determining **31** and advertising **321** respective $CIDAR_c^n$ metrics); receive **322** respective continuously integrated data rate $CIDAR_c^{b_l(n)}$metrics from the one or more neighbor network nodes 1, $b_l(n)$ respectively in association with an identifier of the respective neighbor network node 1, $b_l(n)$; calculate **323** differential metrics $\Delta R_{l,c}^n = CIDAR_c^n - CIDAR_c^{b_l(n)}$ of the determined respective continuously integrated data rate $CIDAR_c^n$ metrics of the network node 1, n and the received respective continuously integrated data rate $CIDAR_c^{b_l(n)}$ metrics of the respective neighbor network node 1, $b_l(n)$; and populate **324** a preference table PFT of the network node 1, n with the respective potential forwarding rules $\langle b_l(n), \Delta R_{l,c}^n \rangle$ for the plurality of communication relations c, the respective potential forwarding rule $\langle b_l(n), \Delta R_{l,c}^n \rangle$ comprising an identifier of a respective neighbor network node 1, $b_l(n)$ and the calculated differential metric $\Delta R_{l,c}^n$ associated with the respective neighbor network node 1, $b_l(n)$.

In other words, the continuously integrated data rate $CIDAR_c^n$ metrics determined by neighbor network nodes 1, $b_l(n)$ are reported to (i.e., advertised to and received by) their neighbors 1, n who store the metric in combination with the next-hop information $b_l(n)$.

The processor 11 may further be configured to, so as to advertise **321** the determined respective continuously integrated data rate $CIDAR_c^n$ metrics: advertise the determined respective continuously integrated data rate $CIDAR_c^n$ metrics via an interior gateway protocol or an exterior gateway protocol.

For example, BGP extensions can be used for carrying CIDAR metrics, but also other reporting protocols / means such as the Internet Control Message Protocol (ICMP), for example.

Given the reporting of CIDAR metrics has tentatively taken place, it would be desirable to distribute requests from more than one client in a manner that compute/network capabilities are best used. As will be explained next, an answer may be found by appropriate TE on the network layer.

As mentioned previously, TE may herein be used to refer to techniques for mapping the given plurality of communication relations c to the given network topology ($n \in \mathcal{N}$, $l \in \mathcal{L}$) with the objective to enhance one or more performance metrics of the network, such as network throughput, for example.

The TE problem of selection of communication relations c for throughput optimization may be understood as a generic optimization problem under an objective function:

$$\max_{x_{l,c}^n \in \{0,1\}} F\left(x_{l,c}^n\right) \qquad \text{(optimization problem)}$$

$$F\left(x_{l,c}^n\right) = \sum_{n \in \mathcal{N}} \sum_{l \in \mathcal{L}^n} \frac{\mathcal{B}_l - CBR_l^n}{\sum_{c \in \mathcal{C}^n} x_{l,c}^n} \cdot \sum_{c \in \mathcal{C}^n} \Delta R_{l,c}^n \qquad \text{(objective function)}$$

$$s.t. \sum_{l \in \mathcal{L}^n} x_{l,c}^n \leq 1, \ \forall c \in \mathcal{C}^n, \ \forall n \in \mathcal{N}$$ (exclusivity constraints on decision variables $x_{l,c}^n$ associated with the set of network links $\mathcal{L}^n$ and the set of communication relations $\mathcal{C}^n$ at the respective network node n, expressing

	the requirement that each communication relation c is assigned to at most one of the set of network links $\mathcal{L}^n$ at the respective network node n)
$\mathcal{N}$	(set of network nodes $\mathcal{N}$ of the network topology)
$\mathcal{L}^n$	(set of network links $\mathcal{L}^n$ at the respective network node n)
$\mathcal{C}^n$	(set of communication relations $\mathcal{C}^n$ at the respective network node n)
$\mathcal{B}_l$	(total bitrate, i.e., capacity, of the respective network link l)
CBR_l^n	(bitrate of constant bitrate (CBR) flows at the respective network node n via the respective network link l)

That is to say, the TE problem amounts to allocating the plurality of communication relations c at the respective network node n to the network links l at the respective network node n in order to maximize the output rate $\frac{\mathcal{B}_l - CBR_l^n}{\sum_{c \in \mathcal{C}^n} x_{l,c}^n}$ per network link l biased by a CIDAR-based weight $\sum_{c \in \mathcal{C}^n} \Delta R_{l,c}^n$.

As can be seen, the network throughput $F\left(x_{l,c}^n\right)$ tends to increase with increasing (sum of) differential metrics $\Delta R_{l,c}^n$.

The optimization problem may be solved by selecting at most one potential forwarding rule $\langle b_l(n), \Delta R_{l,c}^n \rangle$ for the associated communication relation c in order to maximize the given objective function under the given exclusivity constraint. In view of the combinatorial nature of this optimization problem, a greedy algorithm is presented in the following which may yield a near-optimal TE solution in a short time span. The greedy algorithm selects those of the respective potential forwarding rules $\langle b_l(n), \Delta R_{l,c}^n \rangle$ (i.e., which $\langle b_l(n), \Delta R_{l,c}^n \rangle$ tuple, or no tuple) which increase a bandwidth share of the associated relation c on the respective network link l, thereby increasing the network throughput $F\left(x_{l,c}^n\right)$. It should be noted that selecting potential

forwarding rules (*"forward communication relation c from network node n via network link l to neighbor node* $b_l(n)$*"*) implies particular value assignments to the decision variables $x_{l,c}^n$.

As such, the processor 11 may further be configured to, so as to allocate **33** the plurality of communication relations c to the respective network links l in accordance with those of the respective potential forwarding rules $\langle b_l(n), \Delta R_{l,c}^n \rangle$ which increase a bandwidth share of the corresponding associated relation c on the respective network link l: sort **331** the potential forwarding rules $\langle b_l(n), \Delta R_{l,c}^n \rangle$ of the network node 1, n by descending value of the calculated differential metric $\Delta R_{l,c}^n$; mark **332**, in accordance with the descending order of the sorted potential forwarding rules, one of the potential forwarding rules $\langle b_l(n), \Delta R_{l,c}^n \rangle$ for the associated communication relation c for deployment if it improves a throughput $F(\boldsymbol{x})$ of the network; and populate **334** a forwarding table FT of the network node 1, n with the respective entries $\langle c, b_l(n) \rangle$ for the potential forwarding rule $\langle b_l(n), \Delta R_{l,c}^n \rangle$ marked for deployment.

In other words, the sorting **331** takes into account the previous observation that the network throughput $F(x_{l,c}^n)$ tends to increase with increasing (sum of) differential metrics $\Delta R_{l,c}^n$, so that potential forwarding rules $\langle b_l(n), \Delta R_{l,c}^n \rangle$ associated with larger $\Delta R_{l,c}^n$ values should be prioritized, the marking **332** verifies if the respective potential forwarding rule $\langle b_l(n), \Delta R_{l,c}^n \rangle$ actually increases the objective function $F(x_{l,c}^n)$, which represents the network throughput, and the populating **334** incorporates the first potential forwarding rule $\langle b_l(n), \Delta R_{l,c}^n \rangle$ having been verified successfully into the forwarding table FT of the network node 1, n.

Being prepared like this, the respective network node 1, n may perform normal data-plane forwarding. As such, the processor 11 may further be configured to, so as to perform **34** the forwarding for the plurality of communication relations c in accordance with the allocation: receive **341** a packet associated with a particular one of the one or more communication relations c; retrieve **342** the forwarding entry $\langle c, b_l(n) \rangle$ associated with the particular communication relation c from the forwarding table FT; and forward **343** the received packet in accordance with the retrieved forwarding entry $\langle c, b_l(n) \rangle$.

The afore-mentioned plurality of communication relations c may respectively comprise one of: a destination-based communication relation (e.g., IP, TCP/UDP and the like), a flow-based

communication relation (e.g., SDN flow rules), a path-based communication relation (e.g., path-based techniques such as segment routing), a service-based communication relation (e.g., service routing in accordance with a service identifier, such as a hashed uniform resource locator (URL) or a binary structured name in accordance with RFC8609, and possibly a mapping to an egress IP address in accordance with the IETF proposal 'dyncast'), and a label-based communication relation (such as in MPLS).

Especially in case of a service-based communication relation, forwarding of a service request may be realized as extensions within solutions such as (i) constraint-based service routing, (ii) suitable IPv6 header extensions or (iii) solutions such as IP anycast (with service identifier to IP anycast address).

As a result of the distributed TE implementation, the plurality of communication relations c is distributed across the network in a manner that compute/network capabilities are best used.

FIG. 3 illustrates a network node 1, a controller node 2 and corresponding methods 3; 4 of operating the same in accordance with the present disclosure.

The centralized TE implementation may involve a network node and a controller node as follows:

The network node 1, $n \in \mathcal{N}$ is depicted at the top left of FIG. 3 and the controller node 2 is shown at the top right of FIG. 3.

The network node 1, n and the controller node 2 comprise respective processors 11; 21 configured to perform respective methods 3; 4 of operating the respective node 1; 2, which is depicted below the respective node 1; 2 in different levels of detail. For an explanation of the steps **31** and **34** of the method 3, which remain unchanged, reference is made to the corresponding passages provided in connection with FIG. 2 above.

The method 3 of operating the network node 1, n shown at the left of FIG. 3 comprises in its most concise representation the above-identified method steps **31 – 34**, and the network node 1, n comprises the processor 11 and device features corresponding to the method steps **31 – 34**.

Contrary to the distributed TE implementation of FIG. 2 above, the centralized TE implementation provides that the steps **322** – **324** performed by the respective network node 1 are off-loaded to the controller node 2, and that the controller node 2 generates a preference table PFT for the respective network node 1, n comprising the respective potential forwarding rules $\langle b_l(n), \Delta R_{l,c}^n \rangle$.

As such, the processor 11 of the respective network node 1, n may further be configured to, so as to establish **32** the respective potential forwarding rules $\langle b_l(n), \Delta R_{l,c}^n \rangle$ for the plurality of communication relations c based on the respective determined continuously integrated data rate $CIDAR_c^n$ metrics: send **321** the determined respective continuously integrated data rate $CIDAR_c^n$ metrics to a controller node 2 of the network.

The processor 11 may further be configured to, so as to send **321** the determined respective continuously integrated data rate $CIDAR_c^n$ metrics to the controller node 2 of the network: send the determined respective continuously integrated data rate $CIDAR_c^n$ metrics to the controller node 2 in response to a polling indication received from the same.

In other words, the controller node 2 may poll the CIDAR metrics alternatively to periodic and/or occasional reporting by the network nodes 1.

As a result of the off-loading of the establishing **32** functionality, the controller node 2 requires similar functionality with the necessary modifications regarding the centralized execution.

Thus, the method 4 of operating the controller node 2 shown at the right of FIG. 3 comprises in its most concise representation: establishing **42** respective potential forwarding rules $\langle b_l(n), \Delta R_{l,c}^n \rangle$ for a plurality of communication relations c between one or more network nodes 1, n of the network and respective one or more neighbor network nodes 1, $b_l(n)$ connected to the network nodes 1, n via respective network links $l \in \mathcal{L}^n$ based on respective determined continuously integrated data rate $CIDAR_c^n$ metrics. The controller node 1, $b_l(n)$ comprises a processor 21 and device features corresponding to this method step.

Also in the centralized TE implementation, the continuously integrated data rate $CIDAR_c^n$ metrics determined by neighbor network nodes 1, $b_l(n)$ are reported (i.e., advertised and

received), but in such case to the controller node 2 rather than to neighbors 1, n. In return, the controller node 2 may provide a ready-made preference table PFT for the respective network node 1, n comprising the respective potential forwarding rules $\langle b_l(n), \Delta R_{l,c}^n \rangle$.

For this reason, the processor 21 of the controller node 2 may further be configured to, so as to establish **42** the respective potential forwarding rules $\langle b_l(n), \Delta R_{l,c}^n \rangle$ for the plurality of communication relations c between the one or more network nodes 1, n of the network and respective one or more neighbor network nodes 1, $b_l(n)$ connected to the network nodes 1, n via respective network links $l \in \mathcal{L}^n$ based on respective determined continuously integrated data rate $CIDAR_c^n$ metrics: receive **422** respective continuously integrated data rate $CIDAR_c^n$ metrics from the one or more network nodes 1, n respectively in association with an identifier of the respective neighbor network node 1, $b_l(n)$; calculate **423** differential metrics $\Delta R_{l,c}^n = CIDAR_c^n - CIDAR_c^{b_l(n)}$ of the received respective continuously integrated data rate $CIDAR_c^n$ metrics of the one or more network nodes 1, n and the received respective continuously integrated data rate $CIDAR_c^{b_l(n)}$ metrics of the respective one or more neighbor network nodes 1, $b_l(n)$; populate **424** respective preference tables PFT for the one or more network nodes 1, n with the respective potential forwarding rules $\langle b_l(n), \Delta R_{l,c}^n \rangle$ for the respective plurality of communication relations c, the respective potential forwarding rule $\langle b_l(n), \Delta R_{l,c}^n \rangle$ comprising an identifier of a respective one of the one or more neighbor network nodes 1, $b_l(n)$ and the calculated differential metric $\Delta R_{l,c}^n$ associated with the respective one of the one or more neighbor network nodes 1, $b_l(n)$.

Corresponding to similar functionality of the respective network node 1, n, the processor 21 of the controller node 2 may further be configured to, so as to receive **422** the respective continuously integrated data rate $CIDAR_c^n$ metrics from the one or more network nodes 1, n: receive respective continuously integrated data rate $CIDAR_c^n$ metrics from the one or more network nodes 1, n in response to a polling indication sent to the same.

The principle of off-loading the establishing **32** functionality may be applied with regard to the allocating **33** function, too.

That is to say, the centralized TE implementation may provide that the steps **331** – **332** performed by the respective network node 1 are off-loaded to the controller node 2, and that the

controller node 2 provides ready-made entries $\langle c, b_l(n)\rangle$ for the forwarding table FT of the respective network node 1, n for those of the potential forwarding rules $\langle b_l(n), \Delta R_{l,c}^n\rangle$ marked for deployment.

To this end, the processor 21 is further configured to: allocate **43** the plurality of communication relations c to the respective network links l in accordance with those of the respective potential forwarding rules $\langle b_l(n), \Delta R_{l,c}^n\rangle$ which increase a bandwidth share of the associated communication relation c on the respective network link l.

More specifically, the processor 21 may further be configured to, so as to allocate **43** the plurality of communication relations c to the respective network links l in accordance with those of the respective potential forwarding rules $\langle b_l(n), \Delta R_{l,c}^n\rangle$ which increase a bandwidth share of the associated communication relation c on the respective network link l: sort **431** the potential forwarding rules $\langle b_l(n), \Delta R_{l,c}^n\rangle$ of the respective network node 1, n by descending value of the calculated differential metric $\Delta R_{l,c}^n$; mark **432**, in accordance with the descending order of the sorted potential forwarding rules, one of the potential forwarding rules $\langle b_l(n), \Delta R_{l,c}^n\rangle$ for the associated communication relation c for deployment if it improves a throughput $F(\boldsymbol{x})$ of the network; and send **433** respective entries $c, b_l(n)$ for a forwarding table FT of the respective network node 1, n for the potential forwarding rule $\langle b_l(n), \Delta R_{l,c}^n\rangle$ marked for deployment, the respective entry $\langle c, b_l(n)\rangle$ comprising an identifier of the communication relation c associated with the marked potential forwarding rule $\langle b_l(n), \Delta R_{l,c}^n\rangle$ and the neighbor network node 1, $b_l(n)$ of the marked potential forwarding rule $\langle b_l(n), \Delta R_{l,c}^n\rangle$.

In such case, the processor 11 of the respective network node 1 may further be configured to, so as to allocate **33** the plurality of communication relations c to the respective network links l in accordance with those of the respective potential forwarding rules $\langle b_l(n), \Delta R_{l,c}^n\rangle$ which increase a bandwidth share of the corresponding associated relation c on the respective network link l: receive **333** respective entries $\langle c, b_l(n)\rangle$ for the forwarding table FT of the network node 1, n for those of the potential forwarding rules $\langle b_l(n), \Delta R_{l,c}^n\rangle$ marked for deployment, the respective forwarding entry $\langle c, b_l(n)\rangle$ comprising an identifier of the communication relation c associated with the marked potential forwarding rule $\langle b_l(n), \Delta R_{l,c}^n\rangle$ and the neighbor network node 1, $b_l(n)$ of the marked potential forwarding rule $\langle b_l(n), \Delta R_{l,c}^n\rangle$; and populate **334** the

forwarding table FT of the network node 1, n with the respective entries $\langle c, b_l(n)\rangle$ for those of the potential forwarding rules $\langle b_l(n), \Delta R^n_{l,c}\rangle$ marked for deployment.

Since the centralized TE implementation includes substantially the same functionality (or steps) as the distributed TE implementation, the plurality of communication relations c is distributed across the network in a same manner.

The present disclosure has been described in conjunction with various embodiments as examples as well as implementations. However, other variations can be understood and effected by those persons skilled in the art and practicing the claimed matter, from the studies of the drawings, this disclosure and the independent claims. In the claims as well as in the description the word “comprising” does not exclude other elements or steps and the indefinite article “a” or “an” does not exclude a plurality. A single element or other unit may fulfil the functions of several entities or items recited in the claims. The mere fact that certain measures are recited in the mutual different dependent claims does not indicate that a combination of these measures cannot be used in an advantageous implementation. A computer program may be stored/distributed on a suitable medium, such as an optical storage medium or a solid-state medium supplied together with or as part of other hardware, but may also be distributed in other forms, such as via the Internet or other wired or wireless telecommunication systems.

CLAIMS

1. A network node (1) of a network, the network node (1) comprising
a processor (11) configured to:

determine (31) respective continuously integrated data rate metrics for a plurality of communication relations between the network node (1) and one or more neighbor network nodes (1) connected to the network node (1) via respective network links;

establish (32) respective potential forwarding rules for the plurality of communication relations based on the respective determined continuously integrated data rate metrics;

allocate (33) the plurality of communication relations to the respective network links in accordance with those of the respective potential forwarding rules which increase a bandwidth share of the corresponding associated relation on the respective network link; and

perform (34) a forwarding for the plurality of communication relations in accordance with the allocation.

2. The network node (1) of claim 1,
the processor (11) being further configured to, so as to determine (31) the respective continuously integrated data rate metrics for the plurality of communication relations between the network node (1) and the one or more neighbor network nodes (1) connected to the network node (1) via the respective network links:

measure (311) a respective pass-through data rate of the plurality of communication relations at periodic time instants; and

integrate (312) the respective measured data rate over a continuous time period.

3. The network node (1) of claim 1 or claim 2,
the processor (11) being further configured to, so as to establish (32) the respective potential forwarding rules for the plurality of communication relations based on the respective determined continuously integrated data rate metrics:

advertise (321) the determined respective continuously integrated data rate metrics in association with an identifier of the network node (1);

receive (322) respective continuously integrated data rate metrics from the one or more neighbor network nodes (1) respectively in association with an identifier of the respective neighbor network node (1);

calculate (323) differential metrics of the determined respective continuously integrated data rate metrics of the network node (1) and the received respective continuously integrated data rate metrics of the respective neighbor network node (1); and

populate (324) a preference table of the network node (1) with the respective potential forwarding rules for the plurality of communication relations, the respective potential forwarding rule comprising an identifier of a respective neighbor network node (1) and the calculated differential metric associated with the respective neighbor network node (1).

4. The network node (1) of claim 3,
the processor (11) being further configured to, so as to advertise (321) the determined respective continuously integrated data rate metrics:

advertise the determined respective continuously integrated data rate metrics via an interior gateway protocol or an exterior gateway protocol.

5. The network node (1) of claim 1 or claim 2,
the processor (11) being further configured to, so as to establish (32) the respective potential forwarding rules for the plurality of communication relations based on the respective determined continuously integrated data rate metrics:

send (321) the determined respective continuously integrated data rate metrics to a controller node (2) of the network.

6. The network node (1) of claim 5,
the processor (11) being further configured to, so as to send (321) the determined respective continuously integrated data rate metrics to the controller node (2) of the network:

send the determined respective continuously integrated data rate metrics to the controller node (2) in response to a polling indication received from the same.

7. The network node (1) of any one of the claims 3 to 6,
the processor (11) being further configured to, so as to allocate (33) the plurality of communication relations to the respective network links in accordance with those of the respective potential forwarding rules which increase a bandwidth share of the corresponding associated relation on the respective network link:

sort (331) the potential forwarding rules of the network node (1) by descending value of the calculated differential metric;

mark (332), in accordance with the descending order of the sorted potential forwarding rules, one of the potential forwarding rules for the associated communication relation for deployment if it improves a throughput of the network; and

populate (334) a forwarding table of the network node (1) with the respective entries for the potential forwarding rule marked for deployment.

8. The network node (1) of any one of the claims 3 to 6,
the processor (11) being further configured to, so as to allocate (33) the plurality of communication relations to the respective network links in accordance with those of the respective potential forwarding rules which increase a bandwidth share of the corresponding associated relation on the respective network link:

receive (333) respective entries for the forwarding table of the network node (1) for those of the potential forwarding rules marked for deployment, the respective forwarding entry comprising an identifier of the communication relation associated with the marked potential forwarding rule and the neighbor network node (1) of the marked potential forwarding rule; and

populate (334) the forwarding table of the network node (1) with the respective entries for those of the potential forwarding rules marked for deployment.

9. The network node (1) of any one of the claims 1 to 8,
the processor (11) being further configured to, so as to perform (34) the forwarding for the plurality of communication relations in accordance with the allocation:

receive (341) a packet associated with a particular one of the one or more communication relations;

retrieve (342) the forwarding entry associated with the particular communication relation from the forwarding table; and

forward (343) the received packet in accordance with the retrieved forwarding entry.

10. A controller node (2) of a network, comprising
a processor (21) configured to:

establish (42) respective potential forwarding rules for a plurality of communication relations between one or more network nodes (1) of the network and respective one or more neighbor network nodes (1) connected to the network nodes (1) via respective network links based on respective determined continuously integrated data rate metrics; and

allocate (43) the plurality of communication relations to the respective network links in accordance with those of the respective potential forwarding rules which increase a bandwidth share of the associated communication relation on the respective network link.

11. The controller node (2) of claim 10,
the processor (21) being further configured to, so as to establish (42) the respective potential forwarding rules for the plurality of communication relations between the one or more network nodes (1) of the network and respective one or more neighbor network nodes (1) connected to the network nodes (1) via respective network links based on respective determined continuously integrated data rate metrics:

receive (422) respective continuously integrated data rate metrics from the one or more network nodes (1) respectively in association with an identifier of the respective neighbor network node (1);

calculate (423) differential metrics of the received respective continuously integrated data rate metrics of the one or more network nodes (1) and the received respective continuously integrated data rate metrics of the respective one or more neighbor network nodes (1); and

populate (424) respective preference tables for the one or more network nodes (1) with the respective potential forwarding rules for the respective plurality of communication relations, the respective potential forwarding rule comprising an identifier of a respective one of the one or more neighbor network nodes (1) and the calculated differential metric associated with the respective one of the one or more neighbor network nodes (1).

12. The controller node (2) of claim 11,
the processor (21) being further configured to, so as to receive (422) the respective continuously integrated data rate metrics from the one or more network nodes (1):

receive respective continuously integrated data rate metrics from the one or more network nodes (1) in response to a polling indication sent to the same.

13. The controller node (2) of claim 12,
the processor (21) being further configured to, so as to allocate (43) the plurality of communication relations to the respective network links in accordance with those of the respective potential forwarding rules which increase a bandwidth share of the associated communication relation on the respective network link:

sort (431) the potential forwarding rules of the respective network node (1) by descending value of the calculated differential metric;

mark (432), in accordance with the descending order of the sorted potential forwarding rules, one of the potential forwarding rules for the associated communication relation for deployment if it improves a throughput of the network; and

send (433) respective entries for a forwarding table of the respective network node (1) for the potential forwarding rule marked for deployment, the respective entry comprising an identifier of the communication relation associated with the marked potential forwarding rule and the neighbor network node (1) of the marked potential forwarding rule.

14. The network node (1) of any one of the claims 1 to 9 or the controller node (2) of any one of the claims 10 to 13,

the plurality of communication relations respectively comprising one of:

a destination-based communication relation,

a flow-based communication relation,

a path-based communication relation,

a service-based communication relation, and

a label-based communication relation.

15. A method (3) of operating a network node (1) of a network, the method (3) comprising

determining (31) respective continuously integrated data rate metrics for a plurality of communication relations between the network node (1) and one or more neighbor network nodes (1) connected to the network node (1) via respective network links;

establishing (32) respective potential forwarding rules for the plurality of communication relations based on the respective determined continuously integrated data rate metrics;

allocating (33) the plurality of communication relations to the respective network links in accordance with those of the respective potential forwarding rules which increase a bandwidth share of the corresponding associated relation on the respective network link; and

performing (34) a forwarding for the plurality of communication relations in accordance with the allocation.

16. The method (3) of claim 15,

being performed by the network node (1) of any one of the claims 1 to 9 or 14.

17. A method (4) of operating a controller node (2) of a network, the method (4) comprising

establishing (42) respective potential forwarding rules for a plurality of communication relations between one or more network nodes (1) of the network and respective one or more neighbor network nodes (1) connected to the network nodes (1) via respective network links based on respective determined continuously integrated data rate metrics.

18. The method (4) of claim 17,

being performed by the controller node (2) of any one of the claims 10 to 14.

19. A computer program, comprising

executable instructions which, when executed by a processor (11; 21), cause the processor (11; 21) to perform the method (3; 4) of any one of the claims 15 to 18.

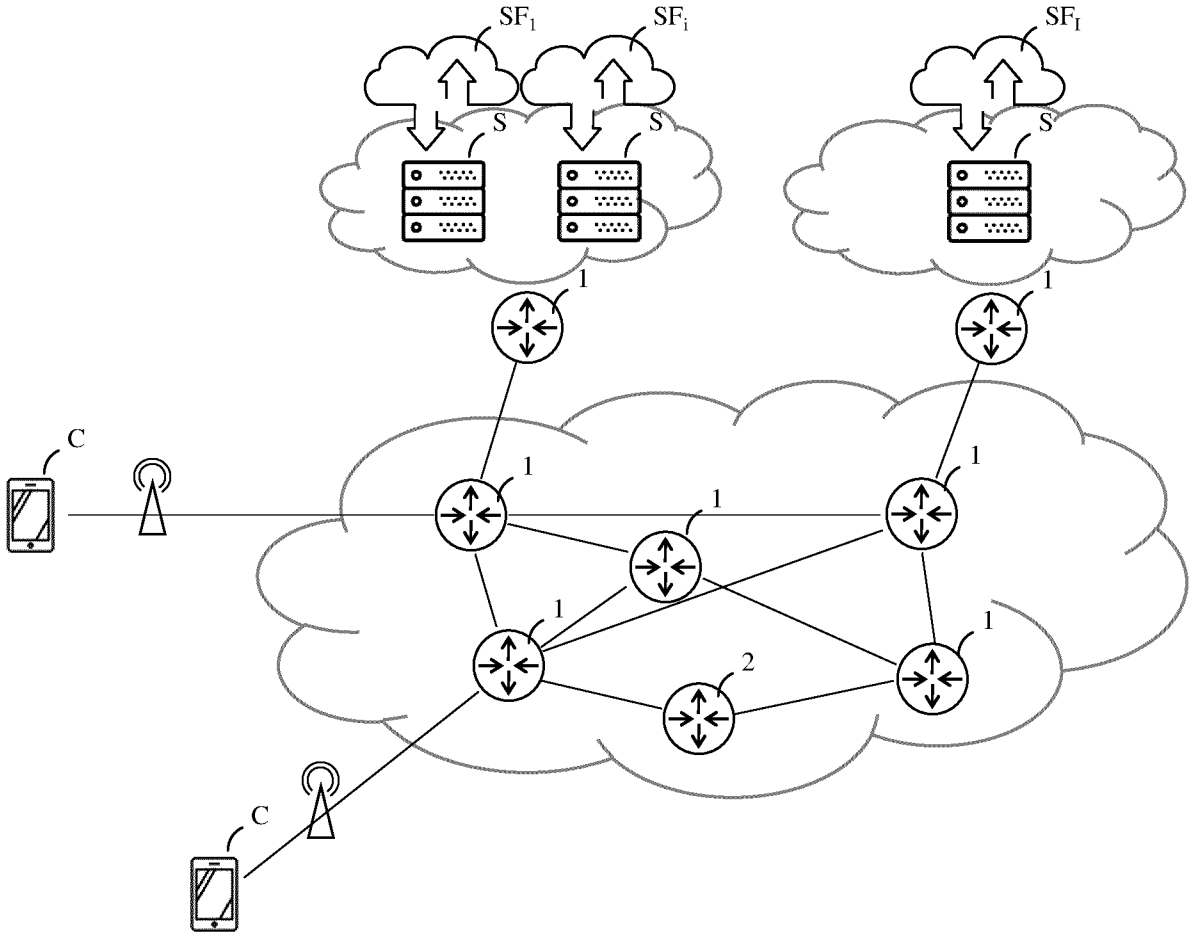


FIG. 1

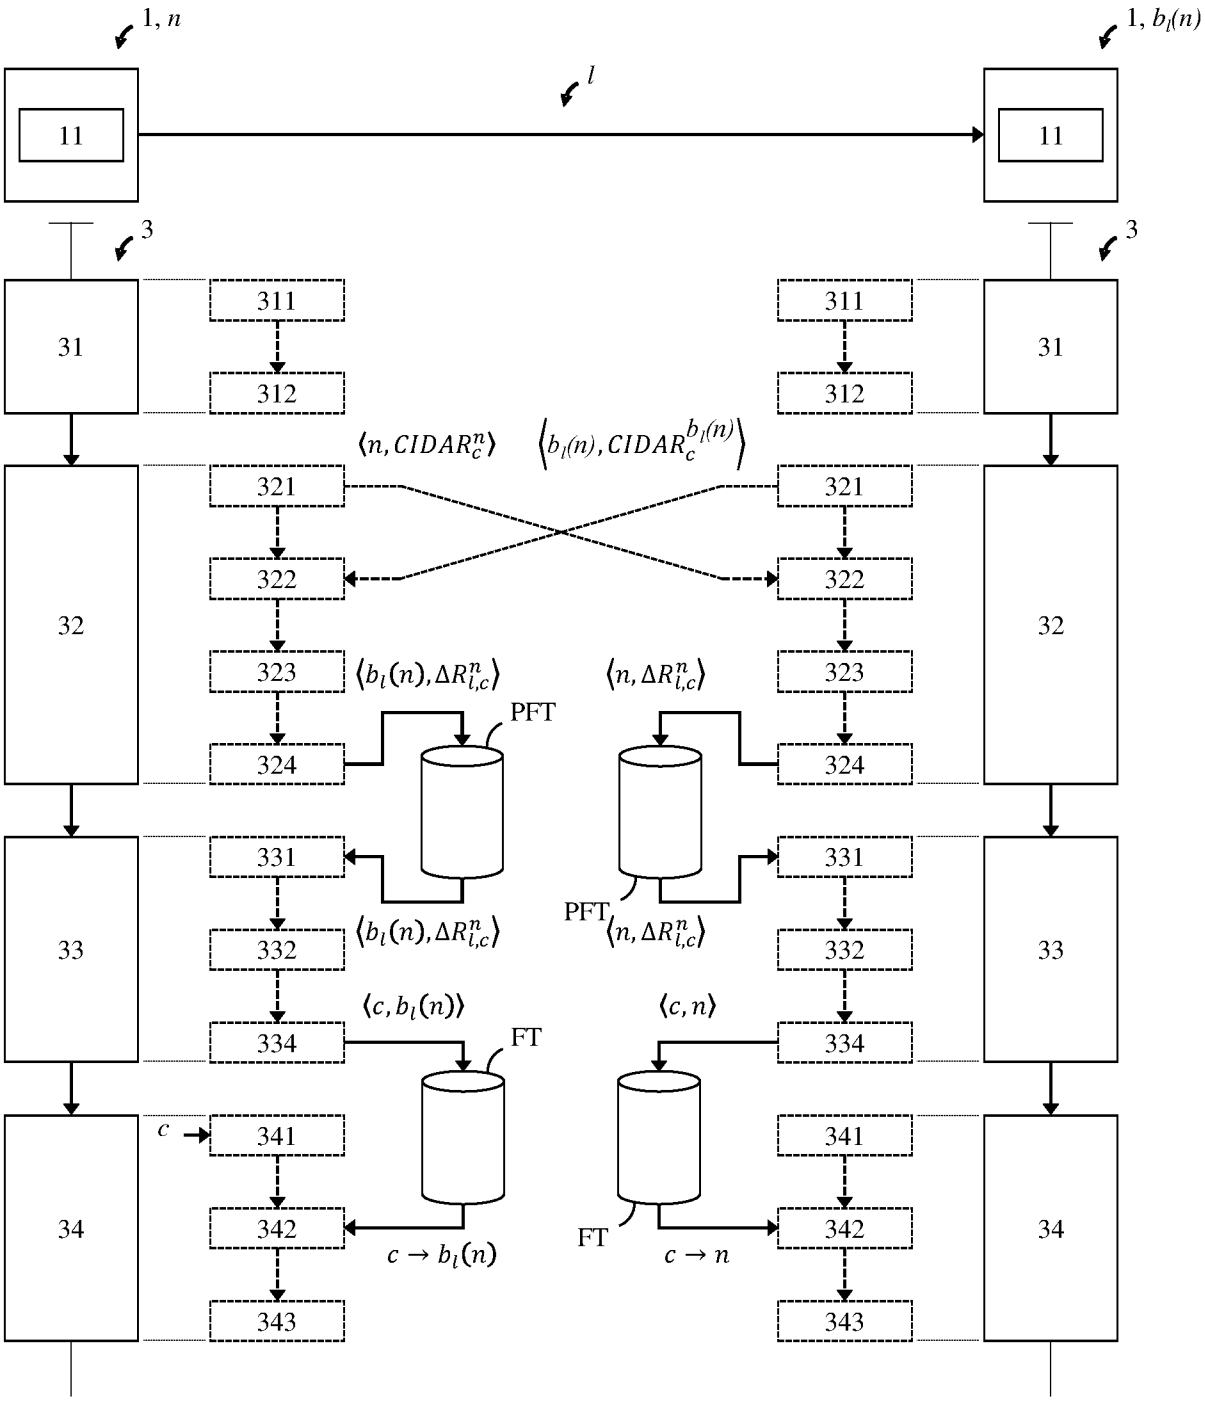

1, n
1, $b_l(n)$
l
11
11
3
3
31
31
311
312
311
312
$\langle n, CIDAR_c^n \rangle$
$\langle b_l(n), CIDAR_c^{b_l(n)} \rangle$
321
322
321
322
32
32
323
324
$\langle b_l(n), \Delta R_{l,c}^n \rangle$
PFT
$\langle n, \Delta R_{l,c}^n \rangle$
323
324
331
332
$\langle b_l(n), \Delta R_{l,c}^n \rangle$
PFT
$\langle n, \Delta R_{l,c}^n \rangle$
331
332
33
33
$\langle c, b_l(n) \rangle$
334
FT
$\langle c, n \rangle$
334
c
341
342
343
$c \rightarrow b_l(n)$
FT
$c \rightarrow n$
341
342
343
34
34

FIG. 2

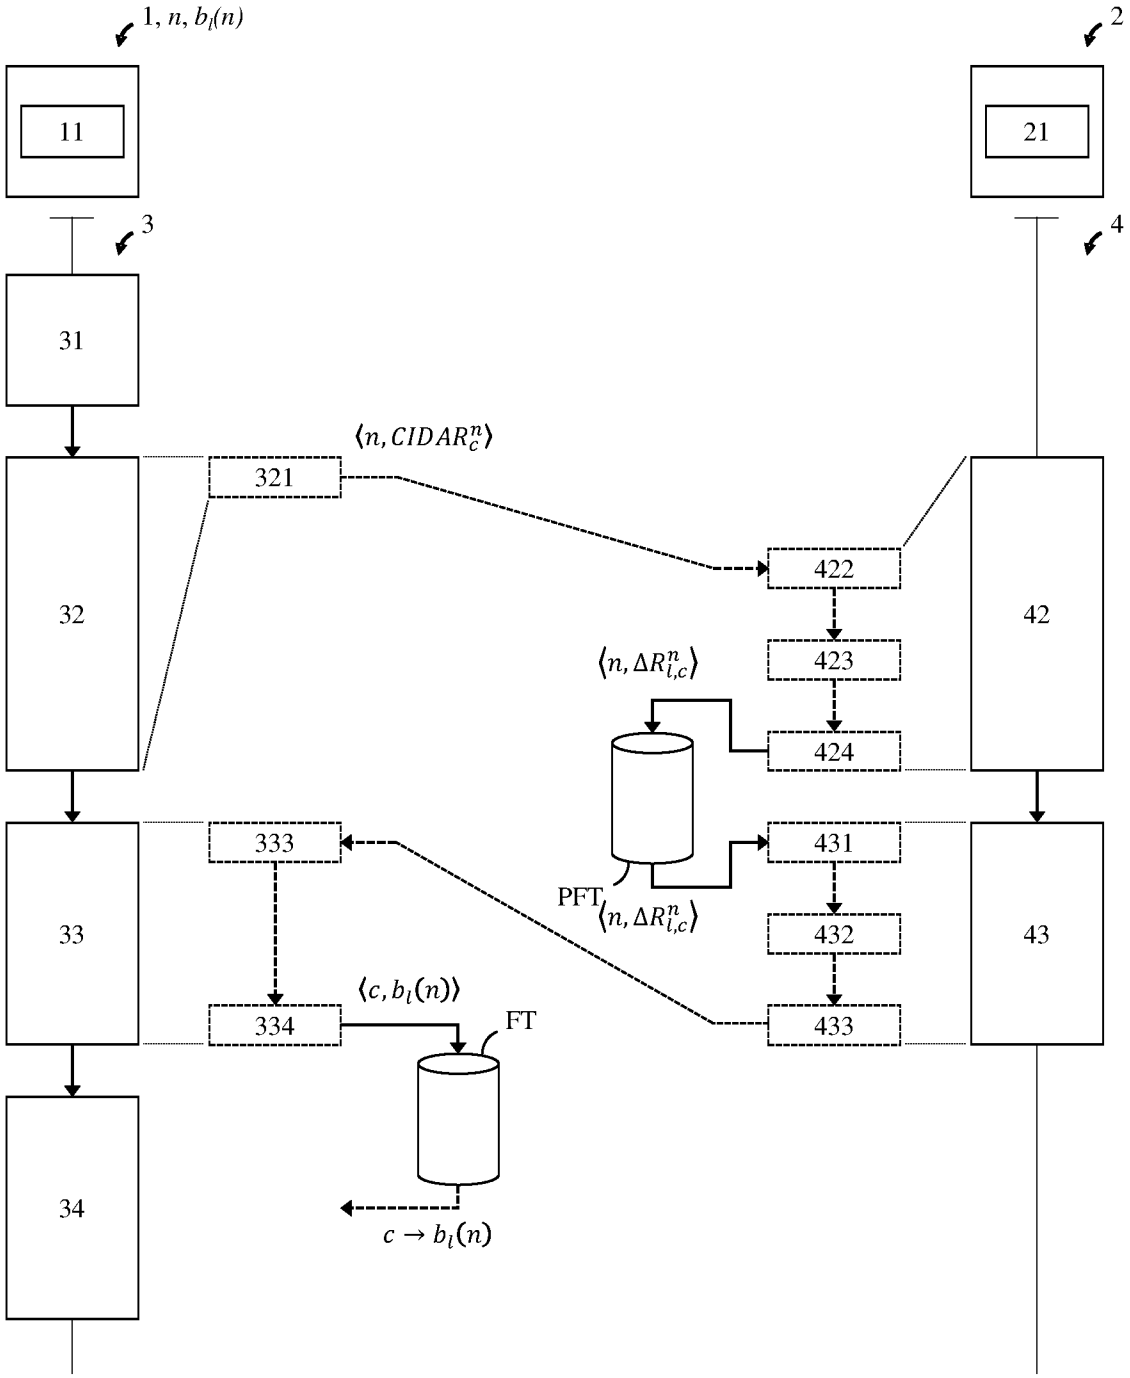

1, n, $b_l(n)$
2
11
21
3
4
31
$\langle n, CIDAR_c^n \rangle$
321
32
422
42
$\langle n, \Delta R_{l,c}^n \rangle$
423
424
333
431
PFT
$\langle n, \Delta R_{l,c}^n \rangle$
33
432
43
$\langle c, b_l(n) \rangle$
334
FT
433
34
$c \rightarrow b_l(n)$

FIG. 3

INTERNATIONAL SEARCH REPORT

International application No
PCT/EP2021/070769

A. CLASSIFICATION OF SUBJECT MATTER
INV. H04L45/24 H04L45/125 H04L45/64
ADD.

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

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
H04L H04W

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

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

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2016/261507 A1 (KWAK JI YOUNG [KR] ET AL) 8 September 2016 (2016-09-08) paragraph [0035] - paragraph [0037]; figure 1 paragraph [0041] - paragraph [0042] paragraph [0052] - paragraph [0055] paragraph [0063] paragraph [0064] paragraph [0069] -----	1-19
X	US 2009/310485 A1 (AVERI JOHN EARNEST [US] ET AL) 17 December 2009 (2009-12-17) paragraph [0053] - paragraph [0060] paragraph [0143] paragraph [0107] - paragraph [0111] -----	1-19

☐ Further documents are listed in the continuation of Box C.

☒ See patent family annex.

* Special categories of cited documents :

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

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

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

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

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

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

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

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

"&" document member of the same patent family

Date of the actual completion of the international search	Date of mailing of the international search report
7 April 2022	20/04/2022

Name and mailing address of the ISA/	Authorized officer
European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Fax: (+31-70) 340-3016	Ciurel, Cristian

1

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No
PCT/EP2021/070769

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2016261507 A1	08-09-2016	KR 20160107825 A	19-09-2016
		US 2016261507 A1	08-09-2016
US 2009310485 A1	17-12-2009	US 2009310485 A1	17-12-2009
		US 2012117273 A1	10-05-2012
		US 2012314578 A1	13-12-2012
		US 2014185445 A1	03-07-2014
		US 2016006658 A1	07-01-2016
		US 2017339059 A1	23-11-2017