

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

(19) World Intellectual Property
Organization
International Bureau



(43) International Publication Date
14 May 2021 (14.05.2021)

(10) International Publication Number
WO 2021/091442 A1

(51) International Patent Classification:

H04W 72/10 (2009.01) *H04W 72/14* (2009.01)
G06N 20/00 (2019.01)

(21) International Application Number:

PCT/SE2019/051124

(22) International Filing Date:

07 November 2019 (07.11.2019)

(25) Filing Language:

English

(26) Publication Language:

English

(71) Applicant: **TELEFONAKTIEBOLAGET LM ERICSSON (PUBL)** [SE/SE]; 164 83 Stockholm (SE).

(72) Inventors: **CHALLITA, Ursula**; Evenemangsgatan 30, 169 56 Solna (SE). **HILTUNEN, Kimmo**; Renbackavägen 2B 32, 02750 Esbo (FI).

(74) Agent: **ERICSSON AB**; Patent Development, Torshamnsgatan 21-23, 164 80 Stockholm (SE).

(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, 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, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States (*unless otherwise indicated, for every kind of regional protection available*): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, 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).

(54) Title: SCHEDULING ENHANCED MOBILE BROADBAND TRANSMISSIONS

300

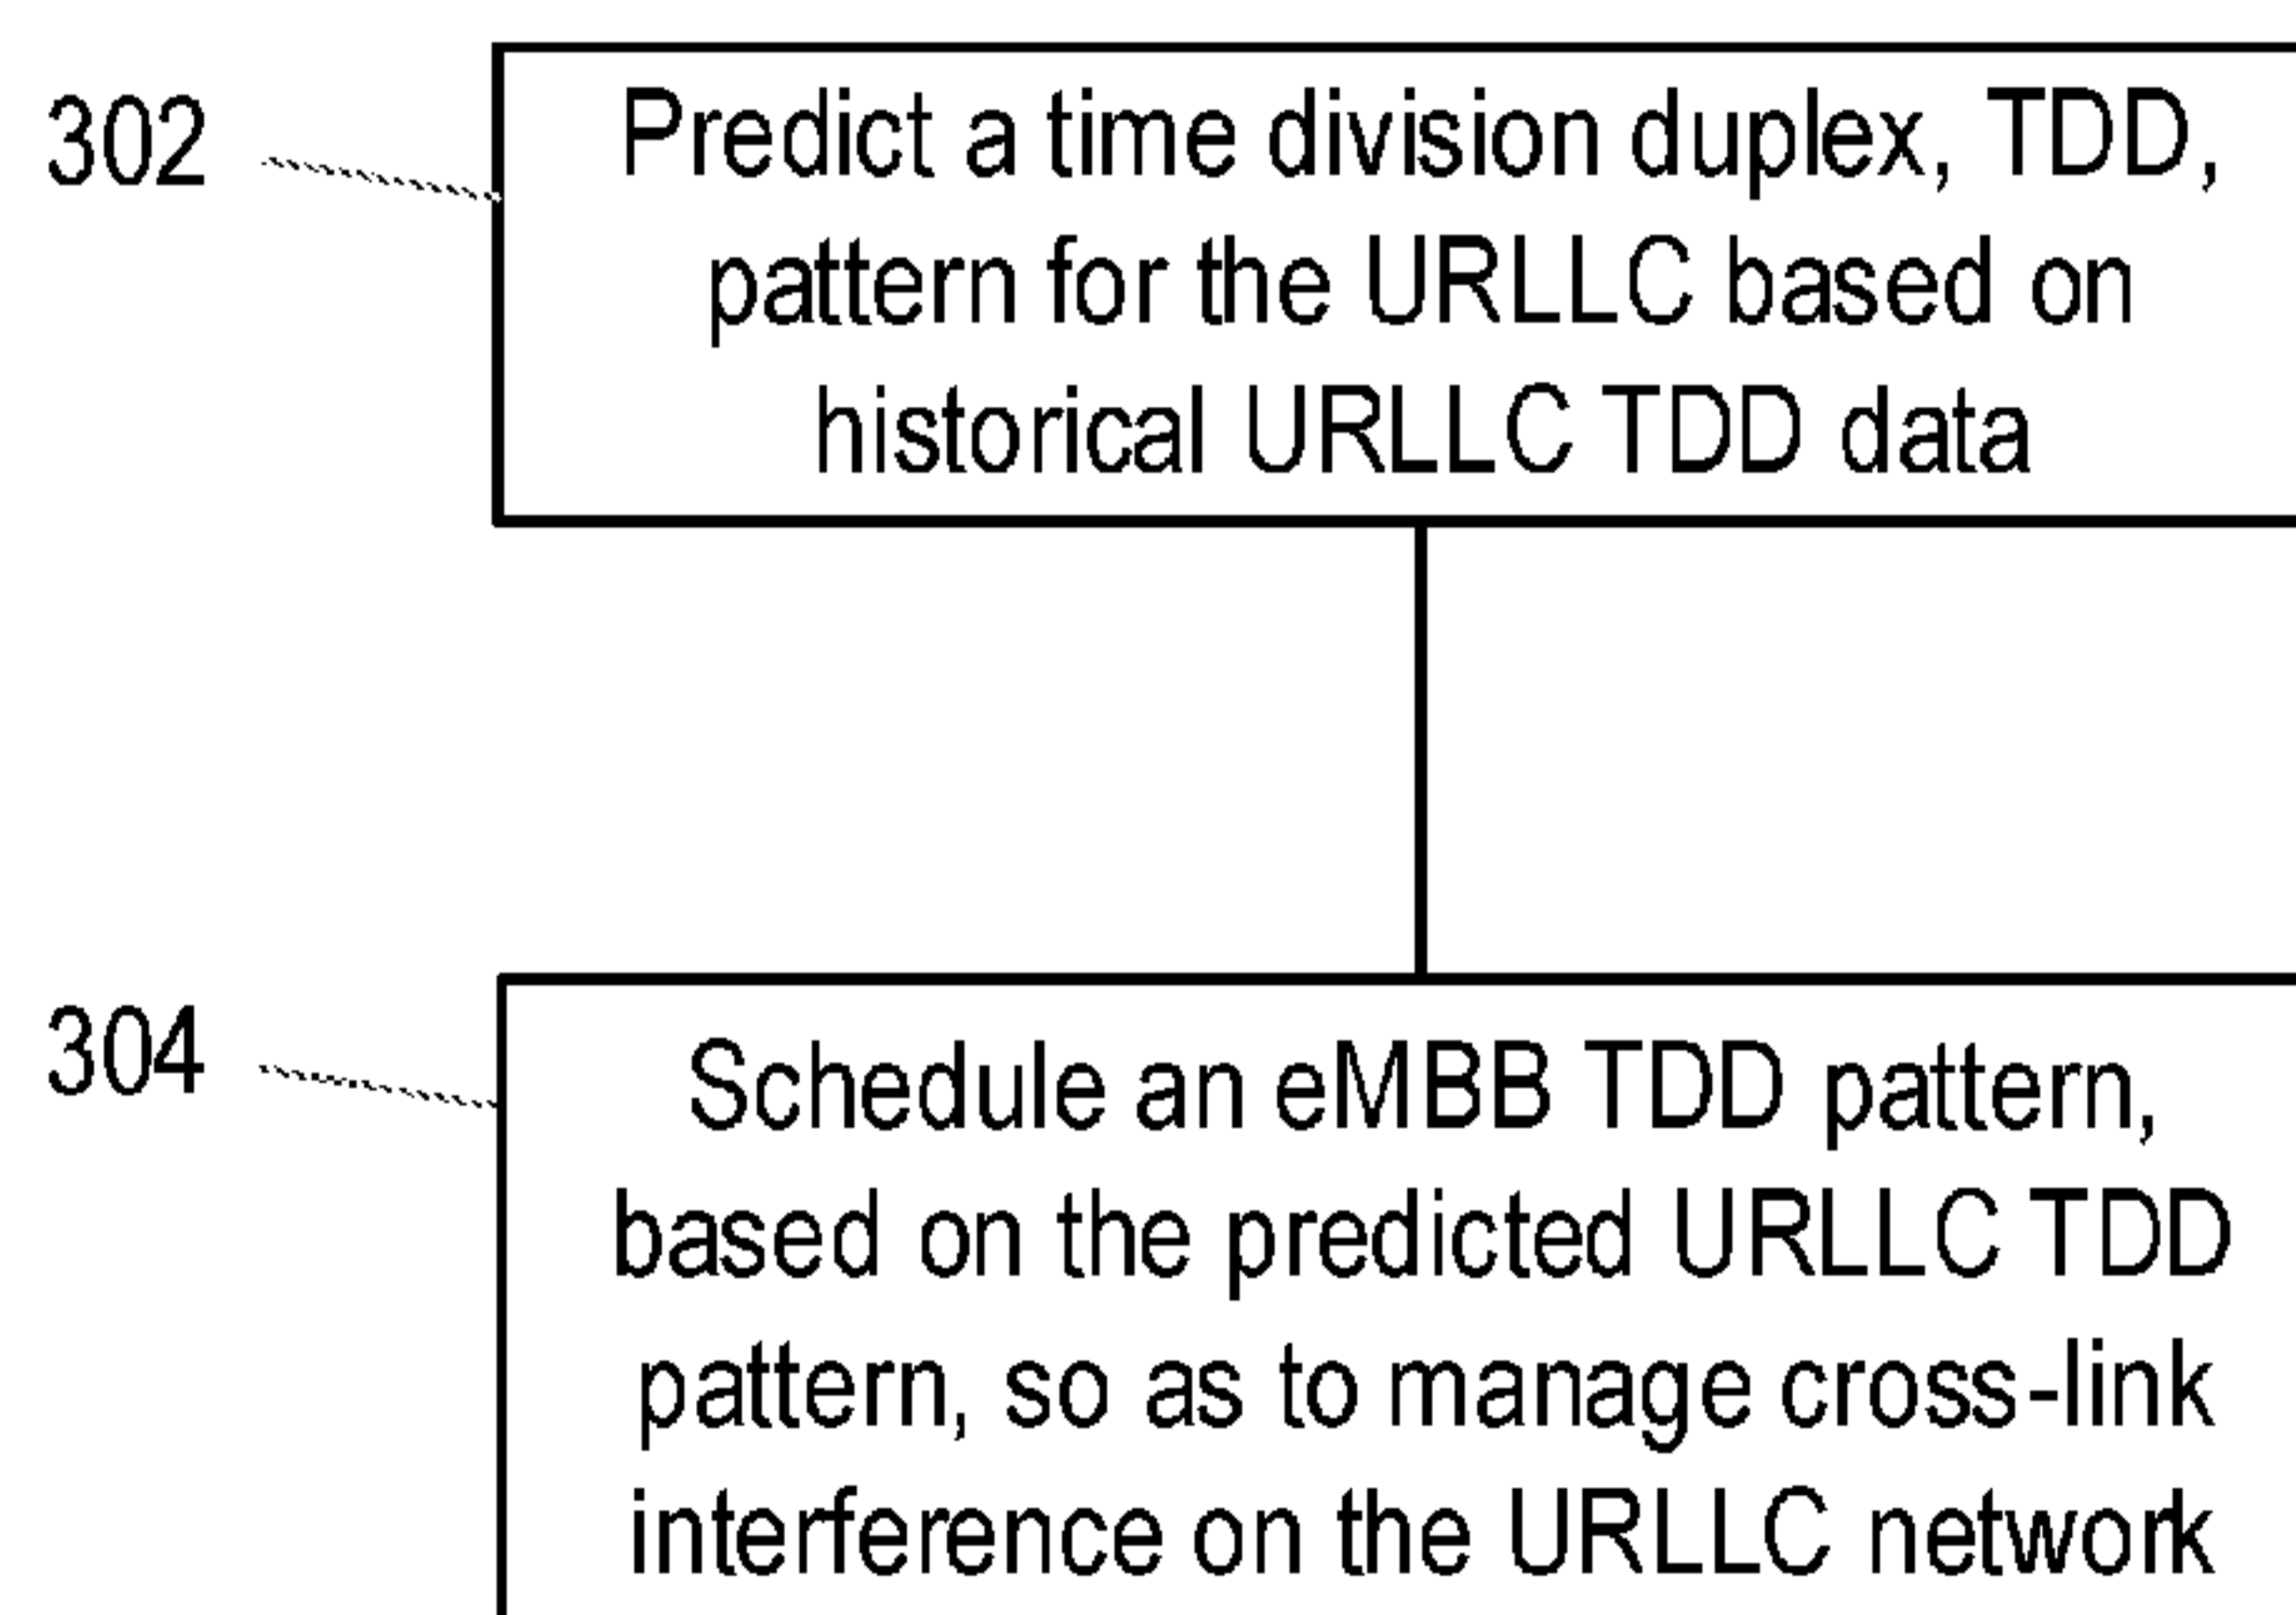


Fig. 3

(57) Abstract: In a method in a node of a communications network for scheduling enhanced mobile-broadband, eMBB, transmissions, the eMBB coverage area interferes with an ultra-reliable low latency communication network, URLLC. In a first block (302) a time division duplex, TDD, pattern for the URLLC is predicted based on historical URLLC TDD data. In a second block (304) an eMBB TDD pattern is scheduled, based on the predicted URLLC TDD pattern, so as to manage cross-link interference on the URLLC network.

[Continued on next page]

WO 2021/091442 A1

Published:

- *with international search report (Art. 21(3))*
- *in black and white; the international application as filed contained color or greyscale and is available for download from PATENTSCOPE*

SCHEDULING ENHANCED MOBILE BROADBAND TRANSMISSIONS

Technical Field

5 This disclosure relates to methods, nodes and systems in a communications network for scheduling enhanced mobile-broadband (eMBB) transmissions. More particularly but non-exclusively, the disclosure relates to scheduling eMBB transmissions so as to manage cross-link interference on an ultra-reliable low latency communication URLLC network.

10

Background

 Fifth generation (5G) mobile networks are envisioned to feature different service classes including ultra-reliable low-latency communications (URLLC), massive machine type communications (mMTC), and enhanced mobile broadband (eMBB). eMBB
15 aims at high spectral efficiency, whilst for URLLC applications, hard latency (e.g., 1 ms) and reliability requirements (target Block Error Rates (BLER) between 10^{-5} and 10^{-9}) are specified. URLLC and eMBB networks are described in the paper by P. Popovski *et al.* (2018) entitled: "5G Wireless Network Slicing for eMBB, URLLC, and mMTC: A
Communication-Theoretic View", IEEE Access, vol. 6, pp. 55765 - 55779, Sept. 2018.

20

 The stringent latency and reliability requirements of URLLC enable new emerging use cases such as factory automation, drone communication, remote control, and remote surgery. An important use case for URLLC is factory automation which may have, for example, required latency of less than 1ms and reliability requirements of 99.999%.

25 Summary

 For 5G use cases such as factory automation (or any other "critical" application), the local URLLC network (e.g. factory) may be overlain by a macro network offering wide area coverage in the same, or neighboring, frequency band. Such macro networks may interfere with the local network, reducing the local network's ability to meet the
30 URLLC latency and reliability requirements.

 In particular, in some scenarios of unsynchronized time division duplex (TDD), the macro and the local factory network may follow different TDD patterns (e.g. different patterns of uplink and downlink transmissions as described in more detail below). For example, the macro network may follow an eMBB-optimized TDD pattern while the local
35 factory network follows an URLLC-optimized TDD pattern. This can mean the eMBB network transmits in the uplink/downlink whilst the URLLC network is transmitting in the downlink/uplink respectively.

As such, base stations in the URLLC network may intermittently experience interference from transmissions from neighboring high-power macro base stations forming part of the eMBB network and having over-lapping coverage areas. This is known as cross-link interference between the base stations, and can have a serious negative impact on the uplink URLLC service availability. More precisely, the latency and reliability requirements of the URLLC users may not be met at different locations within the factory, thus limiting the performance of the factory automation.

It is an object of embodiments herein to address some of these issues. Thus according to a first embodiment herein, there is a method in a node of a communications network for scheduling enhanced mobile-broadband, eMBB, transmissions in scenarios where the eMBB coverage area overlaps with an ultra-reliable low latency communication (URLLC) network. The method comprises predicting a time division duplex, TDD, pattern for the URLLC based on historical URLLC TDD data, and scheduling an eMBB TDD pattern, based on the predicted URLLC TDD pattern, so as to manage cross-link interference on the URLLC network.

As will be described in more detail below, in some embodiments the method may further comprise predicting a traffic load for the eMBB based on historical eMBB traffic load data. The step of scheduling may thus comprise scheduling the eMBB TDD pattern based on the predicted URLLC TDD pattern and the predicted eMBB traffic load.

By using historical data to predict a future TDD pattern for the URLLC, the eMBB node may schedule its transmissions to minimize cross-link interference and thus enable the URLLC network to meet its latency and reliability requirements. URLLC network reliability is essential for applications such as automated factories and/or automated or remote surgery, enabling these technologies to be implemented in a safe and reliable manner.

In some embodiments the step of predicting a time division duplex, TDD, pattern for the URLLC based on historical URLLC TDD data may comprise predicting, at a time, t , the time division duplex, TDD, pattern for the URLLC over a future time interval, $(t;t+s)$, wherein s comprises the duration of the future time interval, and wherein the prediction is based on historical URLLC TDD data obtained in a past time interval $(t-k;t)$, wherein k comprises the duration of the past time interval.

In some embodiments the step of scheduling an eMBB TDD pattern may comprise: at the time t , allocating future eMBB TDD patterns over the future time interval $(t;t+s)$, based on the predicted URLLC TDD pattern over the future time interval $(t;t+s)$ and the predicted eMBB traffic load over the future time interval $(t;t+s)$.

In some embodiments the step of scheduling may comprise scheduling the eMBB TDD pattern so as to meet a service level requirement associated with the URLLC.

In some embodiments the step of scheduling may comprise: minimising a number of slots in which the eMBB transmits in the downlink whilst the URLLC is predicted to transmit in the uplink; and/or minimising a number of slots in which the eMBB transmits in the uplink whilst the URLLC is predicted to transmit in the downlink.

5 In some embodiments the step of scheduling may comprise: scheduling an eMBB uplink slot during a predicted URLLC uplink slot; or scheduling an eMBB downlink slot during a predicted URLLC downlink slot.

10 In some embodiments the step of scheduling may comprise: refraining from scheduling an eMBB downlink slot during a predicted URLLC uplink slot; or scheduling an eMBB downlink slot during a predicted URLLC uplink slot only if required by the eMBB network to meet a service level requirement associated with the eMBB.

15 In some embodiments the step of scheduling may comprise: scheduling an eMBB uplink slot during a predicted URLLC downlink slot for a subset of devices connected to the eMBB network, wherein the subset of devices are located such that the crosslink interference caused by uplink from the subset of devices on the downlink of the URLLC is below a threshold interference level. In some embodiments the threshold interference level may be set based on the signal strength of the subset of devices as measured in the URLLC network.

20 In some embodiments the step of scheduling may comprise: refraining from scheduling an eMBB slot during a predicted URLLC uplink or downlink slot, based on: an amount of downlink data in the buffer; a downlink latency requirement of the eMBB network; and/or an average bit rate requirements of the eMBB.

25 In some embodiments the node applies beamforming, and the step of scheduling may comprise: scheduling the eMBB TDD pattern of beams overlapping a coverage area of the URLLC with a strength above a threshold signal strength, based on the predicted URLLC TDD pattern.

In some embodiments, the step of scheduling may comprise adjusting a numerology associated with the TDD pattern of the eMBB so as to synchronise at least one slot in the TDD pattern of the eMBB with at least one slot in the TDD pattern of the URLLC.

30 In some embodiments, the step of scheduling may comprise adjusting a sub-slot scheduling scheme of the eMBB so as to synchronise at least one slot or sub-slot in the TDD pattern of the eMBB with at least one slot or sub-slot in the TDD pattern of the URLLC.

35 In some embodiments, the step of scheduling may comprise leaving at least one slot or symbol of the eMBB TDD pattern unutilised so as to manage cross-link interference on a slot or sub-slot of the URLLC.

In some embodiments the step of predicting may comprise: training a first model to predict the URLLC TDD pattern using a machine learning process. In some

embodiments the first model comprises a recurrent neural network. In some embodiments the first model may take as input one or more of: a time of day; a day of the week; a previous URLLC pattern; a URLLC load; and an estimate of a future load and/or TDD pattern provided by the URLLC network.

5 In some embodiments training the first model may comprise training the model on training data comprising i) example inputs and ii) corresponding ground truth URLLC TDD patterns. In some embodiments the method may comprise determining a ground truth URLLC TDD pattern based on interference experienced by the eMBB network from the URLLC network. In some embodiments the URLLC and the eMBB may be operated by
10 different network operators.

In some embodiments the method may further comprise receiving a signal from the URLLC network, the signal comprising an indication of an example input and/or an example corresponding ground truth URLLC TDD pattern. In some embodiments the URLLC network and the eMBB network may be operated by the same network operator.

15 In some embodiments the step of predicting traffic load for the eMBB may comprise training a second model to predict the traffic load for the eMBB using a machine learning process. In some embodiments the second model may comprise a recurrent neural network. In some embodiments the second model may take as input one or more of: a time of day; a day of the week; a historical traffic load of the node associated with the eMBB; and
20 a historical traffic load associated with another node neighbouring the node associated with the eMBB.

According to a second aspect there is a node in a communications network, wherein the node is configured for enhanced mobile-broadband, eMBB, transmission with a coverage area that interferes with an ultra-reliable low latency communication network,
25 URLLC. The node comprises: a memory comprising instruction data representing a set of instructions, and a processor configured to communicate with the memory and to execute the set of instructions. The set of instructions, when executed by the processor, cause the processor to: predict a time division duplex, TDD, pattern for the URLLC based on historical URLLC TDD data; and schedule an eMBB TDD pattern, based on the predicted URLLC TDD
30 pattern, so as to manage cross-link interference on the URLLC network.

In some embodiments the processor may be further caused to: predict traffic load for the eMBB based on historical eMBB traffic load data. In some embodiments being caused to schedule may comprise the processor being caused to schedule the eMBB TDD pattern based on the predicted URLLC TDD pattern and the predicted eMBB traffic load.

35 In some embodiments the node comprises one of: a base station; an evolved NodeB; or a network function in a 5G network.

In some embodiments the node is configured to perform the method of the first aspect.

According to a third aspect there is a computer program product comprising a computer readable medium, the computer readable medium having computer readable code embodied therein, the computer readable code being configured such that, on execution by a suitable computer or processor, the computer or processor is caused to perform the method of the first aspect.

Brief Description of the Drawings

For a better understanding and to show more clearly how embodiments herein may be carried into effect, reference will now be made, by way of example only, to the accompanying drawings, in which:

Fig. 1 illustrates two prior art examples of unsynchronized TDD patterns between an eMBB and a URLLC;

Fig. 2 is a graph illustrating interference experienced by a URLLC network at different locations in relation to an eMBB network;

Fig. 3 illustrates a method according to some embodiments herein;

Fig. 4 illustrates a scheme for dividing up a time horizon according to some embodiments herein;

Fig. 5 illustrates a method according to some embodiments herein; and

Fig. 6 illustrates a node according to some embodiments herein.

Detailed Description

Ultra-reliable low-latency communication (URLLC) networks such as those used in automated factory settings, or for automated surgery or similar, may suffer interference from overlying macro networks. Such cross-link interference may result in the URLLC network being unable to meet its latency or reliability requirements. It is an object of embodiments herein to address such issues.

In the frequency domain, User Equipments (UEs) and/or devices are dynamically multiplexed by the orthogonal frequency division multiple access. As will be familiar to the skilled person, in time division duplex (TDD), transmissions over the communications network are divided into slots which are allocated to uplink or downlink transmissions. TDD may be performed over a repeating cycle of uplink and downlink slots, the pattern of which is referred to as the TDD pattern. Each slot in the TDD pattern comprises 14 symbols, each symbol comprising one or more bits of information. The length or duration of the slot depends then on the chosen numerology. A high numerology, i.e., a short slot duration may be seen as an option to improve latency in 5G NR. As another latency enhancement, a concept of non-slot-based transmissions have been introduced as

well, which is here referred to as sub-slot-based scheduling. In case of sub-slot-based scheduling, short transmissions comprising 2, 4 or 7 symbols for downlink and any number of symbols less than 14 for uplink are possible. These short transmissions can start and end in any symbol within a slot.

5 Fig. 1 shows two examples of pairs of transmissions that illustrate unsynchronized TDD patterns. In these examples, each block (labelled “D” for downlink or “U” for uplink) represents a slot 102. Each sequence of four slots (e.g. DDDU and DUDU for the eMBB and URLLC respectively) represent a TDD pattern 104.

10 In the upper example, the eMBB-optimized TDD pattern is “DDDU” while the URLLC-optimized TDD pattern is “DUDU”. Cross-link interference occurs during the blocks highlighted by reference numeral 106, whereby the eMBB and URLLC are unsynchronized with one performing downlink whilst the other performs uplink.

15 In another example, as shown in the lower example in Fig. 1, the eMBB network may follow a DDDU TDD pattern while a local factory (URLLC network) follows a DUDU TDD pattern. This also results in cross link interference in the blocks 106.

20 Generally, as shown in the scenarios shown in Fig. 1, a high level of uplink interference may be seen on the URLLC network when the eMBB network is transmitting in the downlink (as highlighted by the boxes in Fig. 1). This in turn can result in performance degradation of the URLLC network leading to a scenario where latency and reliability requirements cannot be achieved and thus applications of the URLLC network (e.g. factory automation remote surgery etc) are no longer feasible.

25 Fig. 2 shows example URLLC service availability for the case of synchronized (where both networks follow a DUDU TDD pattern) and unsynchronized TDD for different factory locations relative to a neighboring macro cell. In the Figure “Low eMBB” refers to a traffic load of 100 Mbps/km² and “High eMBB” refers to a traffic load of 300 Mbps/km². It can be clearly shown that the URLLC service availability decreases in the case of unsynchronized TDD, and as the factory gets closer to a macro base station, e.g. closer to the center of the macro cell. This is the result of high-power eMBB downlink transmissions during the URLLC uplink transmissions. More information on this may be found in the paper
30 by *Challita et al.* entitled: “*Performance Evaluation for the Co-existence of eMBB and URLLC Networks: Synchronized versus Unsynchronized TDD*”; available at:

<https://arxiv.org/ftp/arxiv/papers/1906/1906.00287.pdf>.

35 To address this issue, Fig. 3 illustrates an example method 300 according to some embodiments herein. The method 300 may be performed by a node in a communications network that is for (e.g. configured or suitable for) scheduling enhanced mobile-broadband, eMBB, transmissions. In scenarios where the coverage area of the eMBB cell or beam interferes with an ultra-reliable low latency communication network, URLLC, the

method 300 may be used to manage cross-link interference on the URLLC network. Briefly, in a block 302 the method 300 comprises predicting a time division duplex, TDD, pattern for the URLLC based on historical URLLC TDD data. Then in a block 304 the method comprises scheduling an eMBB TDD pattern, based on the predicted URLLC TDD pattern, so as to
5 manage cross-link interference on the URLLC network.

One of the realisations that led to the proposed methods and systems herein is that eMBB traffic is typically “best-effort”, e.g., flexible or less latency critical compared to URLLC traffic. Therefore, once the URLLC TDD pattern is predicted, the eMBB traffic is flexible so that it may, for example, be buffered or otherwise scheduled to avoid eMBB
10 downlink transmissions being transmitted in URLLC uplink, or eMBB uplink transmission in URLLC downlink. The idea is to avoid scheduling TDD patterns in a manner that would result in cross-link interference as much as possible, e.g., whenever made possible by the load in the eMBB macro cells. This may be used for example, to minimise the impact of the downlink transmission of the macro base stations on the uplink URLLC transmissions of a factory.
15 Thus increasing (preferably optimising or maximising) the URLLC service availability within the factory, e.g., maximize the percentage of locations within the factory floor where the desired quality of service (QoS) can be guaranteed while also meeting the quality of service for eMBB users.

In more detail, generally the communications network (or telecommunications
20 network) may comprise any combination of wired or wireless links. In some embodiments, the communications network may be configured to operate according to specific standards or other types of predefined rules or procedures. Thus, particular embodiments of the communications network may implement communication standards, such as Long Term Evolution (LTE), and/or other suitable 2G, 3G, 4G, or 5G standards; wireless local area
25 network (WLAN) standards, such as the IEEE 802.11 standards; and/or any other appropriate wireless communication standard, such as the Worldwide Interoperability for Microwave Access (WiMax), Bluetooth, Z-Wave and/or ZigBee standards.

Generally, the node may comprise any component or network function (e.g. any hardware or software module) in the communications network suitable for performing the
30 method 300. In some embodiments the node may comprise the node 600 as described with respect to Figure 6 below. Generally, for example, the node may be involved in resource orchestration and scheduling transmissions/TDD patterns for the eMBB network.

Examples of nodes include, but are not limited to, access points (APs) (e.g., radio access points), base stations (BSs) (e.g., radio base stations, Node Bs, evolved Node
35 Bs (eNBs) and NR NodeBs (gNBs)). Further examples of nodes include but are not limited to core network functions such as, for example, core network functions in a Fifth Generation Core network (5GC). Examples of 5GC network functions include, but are not limited to the

Access and Mobility Management function (AMF), Session Management function (SMF) and Network Slice Selection Function (NSSF).

The node may be configured for making eMBB transmissions, e.g. the method 300 may be used by the node to schedule its own eMBB transmissions. Alternatively, the
5 node may be configured to communicate with another node, the other node being configured to perform the scheduled transmissions, e.g. the node may use the method 300 to schedule the eMBB transmissions of another node.

eMBB transmissions may comprise transmissions that comply with the 3GPP New Radio (NR) standard.

10 The URLLC may be associated with, for example, a factory (e.g. a “smart factory”) for performing factory automation; drone communication for navigating drones; automated vehicles; and/or remote surgery applications. The skilled person will appreciate however that these are only examples, and that the URLLC may be associated with any other application.

15 Generally the URLLC network may comply with the 3GPP New Radio (NR) standard.

Generally, transmissions made by the eMBB may interfere with transmissions of a URLLC. As such, the coverage area of the node performing the eMBB transmissions partially or fully overlaps with the coverage area of an ultra-reliable low latency
20 communications network URLLC. In this sense, for example, the eMBB may neighbor the URLLC network.

In some embodiments, the eMBB and URLLC may operate on the same frequency or channel. In other embodiments, the eMBB URLLC may operate on different frequencies. Generally, the URLLC and eMBB may operate on any frequency combination
25 for which interference may be observed by the URLLC. For example, cross link interference is seen in adjacent channels and thus the URLLC and eMBB networks may be co-channel or adjacent channel deployments.

Generally, the eMBB and URLLC networks may be operated by the same network operator or different network operators, as will be discussed in more detail below.

30 Turning to Fig. 4, Fig. 4 illustrates a framework for performing the method 300 according to some embodiments. In this embodiment, the time horizon is divided into time intervals or windows, each of size T . It is noted that a time interval may also be referred to at any point herein as a time window. Each time window is divided into two sub-windows of size k and s , respectively, as depicted in Fig. 4. The first sub-window starts at time instant $(t-k)$
35 and ends at time instant t while the second starts at time instant t and ends at time instant $(t+s)$.

At time t , the method 300 is performed. In such embodiments, over the period k , the node (e.g. macro base station) learns the historical URLLC activity and TDD pattern over time window $(t-k; t)$. The step of predicting a TDD pattern may thus comprise predicting the TDD pattern over the interval $(t;t+s)$.

5 Put another way, in some embodiments the step of predicting 302 a time division duplex, TDD, pattern for the URLLC based on historical URLLC TDD data may comprise predicting, at a time, t , the time division duplex, TDD, pattern for the URLLC over a future time interval, $(t;t+s)$, wherein s comprises the duration of the future time interval. The prediction may be based on historical URLLC TDD data obtained in a past (e.g. previous)
10 time interval $(t-k;t)$, wherein k comprises the duration of the past time interval.

The node may obtain the historical URLLC TDD pattern (e.g. in some embodiments over the time window $(t-k; t)$ as described above) as follows. As a first example, in some embodiments, the node may monitor the URLLC TDD pattern by performing or collecting measurements of interference from devices or user equipment (UEs)
15 connected to the node e.g., by monitoring the interference level from the URLLC network on the macro UEs, the node may infer the TDD pattern of the URLLC network.

During deployment, the eMBB network can infer the URLLC TDD pattern over $(t-k; t)$ based on the interference level on eMBB users surrounding the URLLC network (e.g. factory). For instance, the interference level on eMBB users close to a factory is higher in
20 uplink when the URLLC is transmitting in the downlink compared to when the URLLC is also transmitting in the uplink. Furthermore, the interference level from URLLC on eMBB is proportionally mapped to the resource utilization of the URLLC network. It is noted that it may be desirable that macro cell transmissions be avoided during time periods when such measurements are performed on the URLLC network.

25 This embodiment may be relevant, for example, where the eMBB network and the URLLC network are operated by different network operators or vendors and information sharing between the operators is limited. In this manner the eMBB network may be able to infer a historical TDD pattern without receiving any information from the URLLC network.

As a second example, in some embodiments the URLLC network may signal its
30 TDD pattern directly to the eMBB network. For example the URLLC network may send one or more signals to the macro network, the one or more signals comprising information about the historical load and/or TDD pattern of the URLLC. For example, the URLLC may provide such information about regarding the time interval (or window) $(t-k; t)$.

In some embodiments, the URLLC network may provide, alternatively or
35 additionally, an estimate of the expected load and TDD pattern for the coming time period. As such, the step 302 may comprise obtaining a prediction of a TDD pattern for the URLLC network from the URLLC network itself. Such embodiments may be possible in particular if

the URLLC network belongs to the same operator as the macro network, which is relevant for many co-channel deployments between macro and local URLLC networks. For adjacent channel deployments, multiple operators may be involved and therefore the first example may be more convenient.

5 Based on the historical activity of the URLLC network in the step 302, the method comprises predicting 302 a time division duplex, TDD, pattern for the URLLC based on the historical URLLC TDD data. In some embodiments, this step comprises the node (e.g. macro BS(s)) predicting a future URLLC TDD pattern over a time window of (t; t+s).

10 In some embodiments, the step 302 may comprise predicting a sequence of actions. For example, a sequence of TDD patterns.

15 In some embodiments the prediction is made using a first trained model such as, for example, a model trained using a machine learning process. In some embodiments, the step of predicting may be performed using first model such as a trained neural network. In some examples, the first model may comprise a recurrent neural network (RNN). In one example, the first model may comprise a long short term memory (LSTM) cell. RNN are suitable for time-dependent data. RNN is suitable for the embodiments herein since a TDD pattern is a function of load and load can be characterized by a time series. For instance, the load level at t is a function of the load at t-1, t-2, t-3. There thus exhibits a time dependency.

20 The skilled person will be familiar with machine learning and neural networks, but in brief, neural networks are a type of supervised machine learning model that can be trained to predict a desired output for given input data. Neural networks are trained by providing training data comprising example input data and the corresponding "correct" or ground truth outcome that is desired. Neural networks comprise a plurality of layers of neurons, each neuron representing a mathematical operation that is applied to the input data. The output of each layer in the neural network is fed into the next layer to produce an output. For each piece of training data, weights associated with the neurons are adjusted until the optimal weightings are found that produce predictions for the training examples that reflect the corresponding ground truths.

25 The skilled person will appreciate however that neural networks are merely an example and that any model or prediction scheme that may be used to predict a URLLC TDD scheme could equally be used. Examples of other models include but are not limited to random forest models.

30 Generally, the method 300 may further comprise training a first model to predict the URLLC TDD pattern using a machine learning process. The first model may take as input one or more of: a time of day; a day of the week; a previous URLLC pattern; a URLLC load; and an estimate of a future load and/or TDD pattern provided by the URLLC network. Put another way, these may be input features or input parameters to the first model

(or prediction scheme). The skilled person will appreciate that these are merely examples and that other inputs may also be provided to the first model.

Where the first model takes as input a previous URLLC pattern or a URLLC load, these may comprise the URLLC TDD pattern over the time interval k (e.g. the slots in the time interval $(t-k;t)$ as described above.

In some embodiments, training the first model may comprise training the model on training data comprising i) example inputs and ii) corresponding ground truth URLLC TDD patterns. As an example, in an embodiment where the first model takes as input three input parameters comprising i) Time of Day, ii) "Actual URLLC Load in the time period $(t-k;t)$ and iii) Estimated Future Load in the time period $(t;t+s)$, the training data may be in the following format:

	TimeOfDay	URLLCLoad($t-k;t$)	EstimatedFutureLoad($t;t+s$)	ActualURLLCTDDPATTERN($t;t+s$)
	12:04:48	25, 25, 30, 40, 15	40, 10, 50, 5	DUDU
15	12:56:43	20, 15, 25, 15, 25	10, 40, 10, 5	UDUU

In such an example the field ActualURLLCTDDPATTERN($t;t+s$) represents the ground truth, e.g. the TDD pattern that the URLLC network actually performed in the subsequent time period $(t;t+s)$ for each example. In this example, the first model is thus trained to predict the field: ActualURLLCTDDPATTERN($t;t+s$) from the input parameters: TimeOfDay, URLLCLoad($t-k;t$) and EstimatedFutureLoad($t;t+s$). In this example, the units of the URLLC Load and EstimatedFutureLoad are packets/s/m². $k=5$ and $s=4$ in the above example.

In some embodiments the method may comprise determining a ground truth URLLC TDD pattern based on interference experienced by the eMBB network from the URLLC network (as described above). This may be particularly relevant for compiling a training data set in embodiments where the URLLC and the eMBB networks are operated by different network operators where information sharing may be limited. Alternatively or additionally, the URLLC may directly signal the eMBB network with example inputs and ground truth outputs as described above. The method may therefore further comprise receiving a signal from the URLLC network comprising an indication of an example input and/or an example corresponding ground truth URLLC TDD pattern. This may be particularly relevant if the eMBB and URLLC networks are operated by the same operator and thus direct signalling may be more easily facilitated.

The skilled person will be familiar with methods of training a neural network using training data (e.g. gradient descent etc.) and appreciate that the training data may

comprise many hundreds or thousands of rows of training data (depending on the accuracy required of the trained model), obtained in a diverse range of network conditions.

In some embodiments, the method 300 may further comprise predicting traffic load for the eMBB based on historical eMBB traffic load data. The step of scheduling 304 (as
5 will be described in more detail below) may then comprise scheduling the eMBB TDD pattern based on the predicted URLLC TDD pattern and the predicted eMBB traffic load.

In terms of the schematic in Fig. 4, in some embodiments the step of predicting traffic load for the eMBB based on historical eMBB traffic load data may comprise predicting, at a time, t , traffic load for the eMBB over a future time interval, $(t;t+s)$, wherein s comprises
10 the duration of the future time interval. The prediction may be based on historical traffic load for the eMBB obtained in a past time interval $(t-k;t)$, wherein k comprises the duration of the past time interval. Put another way, based on the historical activity of the eMBB users, each macro base station predicts the future traffic load of its associated users over a time window of $(t; t+s)$.

In some embodiments the prediction of the eMBB load is made using a second
15 trained model such as, for example, a model trained using a machine learning process. In some embodiments, the step of predicting traffic load for the eMBB based on historical eMBB traffic load data may be performed using a second model such as a trained neural network. In some examples, the second model may comprise a recurrent neural network. In one
20 example, the first model may comprise a long short term memory (LSTM) cell.

Neural networks were described above with respect to the first model and the details therein will be understood to apply equally to the second model. The skilled person will further appreciate that the second model may comprise any other type of model or prediction scheme suitable for predicting a future eMBB load based on historical eMBB traffic
25 load information. Examples of other models include but are not limited to random forest models.

Generally, the method 300 may further comprise training the second model to predict the future eMBB load using a machine learning process.

The second model may take as input one or more of: a time of day; a day of the
30 week; a historical traffic load of the node associated with the eMBB; and a historical traffic load associated with another node neighbouring the node associated with the eMBB. Put another way, these may be input features or input parameters to the second model (or prediction scheme). The skilled person will appreciate that these are merely examples and that other inputs may also be provided to the second model.

Where the second model takes as input a historical traffic load of the node
35 associated with the eMBB or a historical traffic load associated with another node neighbouring the node associated with the eMBB, these may comprise historical traffic loads

over the time interval k (e.g. the slots in the time interval $(t-k;t)$) as described above.

In some embodiments, training the second model may comprise training the model on training data comprising i) example inputs and ii) corresponding ground truth eMBB traffic load information. Generally the node, in its capacity as scheduler of the eMBB

5 transmissions, will be able to record its historic traffic load for use as training data.

As an example, in an embodiment where the second model takes as input two input parameters comprising i) Time of Day, and ii) eMBB Load in the time period $(t-k;t)$, the training data may be in the following format:

10	TimeOfDay	eMBBLoad($t-k;t$)	ActualeMBBLoad($t;t+s$)
	12:04:48	100 Mbps/km ²	110 Mbps/km ²
	12:56:43	90 Mbps/km ²	100 Mbps/km ²

In such an example the field $ActualeMBBLoad(t;t+s)$ represents the ground truth, e.g. the eMBB load in the subsequent time period $(t;t+s)$ for each example. In this
 15 example, the second model is thus trained to predict the field: $ActualeMBBLoad(t;t+s)$ from the input parameters: TimeOfDay, and eMBBLoad($t-k;t$). In the example above, $k=1$ and $s=1$. However, it will be appreciated that in other examples, a sequence of values may be provided as input and a sequence of future values may be predicted as output e.g., the input and outputs may be vectors and not scalars.

20 As noted above with respect to the first model, the skilled person will be familiar with methods of training a neural network using training data (e.g. gradient descent etc.) and appreciate that the training data may comprise many hundreds or thousands of rows of training data (depending on the accuracy required of the trained model), obtained in a diverse range of network conditions.

25 The first model (and the second model in embodiments where the second model is generated) may be trained according to data driven schemes during which a large amount of data traffic from the network is collected over multiple days. The node performing the method 300 (e.g. the macro base station) may be trained offline to learn to predict a future URLLC TDD pattern and/or the future eMBB load of its associated users using the first
 30 and/or second models over a time window of $(t; t+s)$ respectively based on the respective considered input features.

The first model and/or second model may be periodically updated by performing further training and providing updates to the first and/or second models. Further training may be triggered, for example, by changes to the URLLC or eMBB network (e.g. new
 35 factory equipment, new macro base stations, changes in beamforming and the like). In this way, the model may be trained to respond to changes in the network conditions and thus the accuracy of the first and/or second models may be ensured over a longer period of time.

Turning back now to the method 300, after block 302 (e.g. after the TDD pattern for the URLLC is predicted), the method 300 moves on to block 304: scheduling an eMBB TDD pattern, based on the predicted URLLC TDD pattern, so as to manage cross-link interference on the URLLC network.

5 In this step, neighboring eMBB nodes to the URLLC (e.g. neighbouring macro base stations) adjust their corresponding TDD patterns over a time window of size s slots based on the predicted URLLC TDD pattern.

In embodiments where the eMBB traffic load is also predicted (as described above), the scheduling may be based on the predicted URLLC TDD pattern and the
10 predicted eMBB traffic load over the time window $(t; t+s)$. For example, at the time the eMBB TDD pattern may be scheduled by allocating future eMBB TDD patterns over the future time interval $(t;t+s)$, based on both the predicted URLLC TDD pattern over the future time interval $(t;t+s)$ and the predicted eMBB traffic load over the future time interval $(t;t+s)$.

For example, the input to this step may comprise the predicted URLLC TDD
15 pattern (from Step 302). In some embodiments, a predicted eMBB traffic load over time window $(t; t+s)$ may also be input. In embodiments where the predicted eMBB traffic load over time window $(t; t+s)$ is provided, this may be used to approximate the number of eMBB downlink TDD slots required for serving the predicted eMBB load over time window $(t; t+s)$.

Generally, the cross-link interference may be managed by the eMBB network
20 controlling the number of slots during which the eMBB is transmitting in the downlink and URLLC is transmitting in the uplink. Thus in some embodiments, the step of scheduling comprises minimising a number of slots in which the eMBB transmits in the downlink whilst the URLLC is predicted to transmit in the uplink; and/or minimising a number of slots in which the eMBB transmits in the uplink whilst the URLLC is predicted to transmit in the downlink. In
25 this way cross-link interference is avoided wherever possible and thus the URLLC latency and reliability requirements may be guaranteed.

More generally, the step 304 of scheduling may comprise scheduling the eMBB TDD pattern so as to meet a service level requirement associated with the URLLC. For example, some eMBB downlink/uplink transmissions may be permitted in URLLC predicted
30 uplink/downlink slots (respectively) so long as the interference caused is not significant enough to prevent the URLLC from meeting the service level requirements required of URLLC networks.

To this end, in some embodiments, the step of scheduling may comprise scheduling an eMBB uplink slot during a predicted URLLC uplink slot; or scheduling an
35 eMBB downlink slot during a predicted URLLC downlink slot. In this way, cross link interference may be avoided in those slots. In some embodiments, cross-link interference may be substantially avoided by only scheduling eMBB uplink slots during predicted URLLC

uplink slots and only scheduling eMBB downlink slots during predicted URLLC downlink slots.

Put another way, step 304 may comprise refraining from scheduling (e.g. not scheduling) eMBB downlink slots during predicted URLLC uplink slots and/or refraining from scheduling (e.g. not scheduling) eMBB uplink slots during predicted URLLC downlink slots.

In the above, it is assumed that the eMBB transmissions are entirely flexible and can thus be buffered wherever necessary in order to give preference to the URLLC transmissions. In some embodiments the step of scheduling may comprise scheduling an eMBB downlink slot during a predicted URLLC uplink slot or scheduling an eMBB uplink slot during a predicted URLLC downlink slot, (only) if required by the eMBB network to meet a service level requirement associated with the eMBB. In other words, in some embodiments, the eMBB may prioritise its traffic if buffering would lead the eMBB to fail to meet its own reliability or quality of service criteria. More generally therefore, the step of scheduling may comprise refraining from scheduling an eMBB slot during a predicted URLLC uplink or downlink slot, based on: an amount of downlink data in the buffer; a downlink latency requirement of the eMBB network; and/or an average bit rate requirements of the eMBB.

In another embodiment, the scheduling may be performed selectively, so that only traffic to devices with signal strength in range of the URLLC network may be scheduled as above, whilst traffic to devices not in range of the URLLC network may be treated normally. In other words, an eMBB uplink slot may be scheduled during a predicted URLLC downlink slot for a subset of devices connected to the eMBB network if the subset of devices are located such that the crosslink interference caused by uplink from the subset of devices on the downlink of the URLLC is below a threshold interference level. The threshold interference level may be set based on the signal strength of the subset of devices as measured in the URLLC network. Put another way, the URLLC network may measure the signal strength of the eMBB devices, and signal to the eMBB network if scheduling restrictions are needed (e.g, if the received signal strength exceeds the threshold interference level). Alternatively or additionally, eMBB devices can measure the signal strength from the URLLC, and report it back to the eMBB network, which may then take scheduling actions based on the threshold interference level. In some examples, the threshold interference level may be set so as to ensure a latency, reliability or other QoS criteria may be met by the URLLC network. The threshold interference level may also have been agreed between the eMBB network and the URLLC network.

To summarise, in a particular embodiment of the method 300, if the predicted URLLC TDD pattern is downlink for the next time slot, the eMBB transmits in the downlink if there is downlink data in the buffer. Alternatively, if there is no downlink data in the buffer, or if the downlink latency and the average bit rate requirements allow it, the eMBB network can

either be silent, or schedule uplink transmissions from users located sufficiently far away from the URLLC network. The sub-group of users, which can be considered for uplink transmissions can for example be decided based on the signal strength received by the macro users from the URLLC network. In this embodiment, if the predicted URLLC TDD
5 pattern is downlink for the next time slot, eMBB transmits in the downlink if there is downlink data in the buffer. Alternatively, if there is no downlink data in the buffer, or if the downlink latency and the average bit rate requirements allow it, the eMBB network can either be silent, or schedule uplink transmissions from users located sufficiently far away from the URLLC network. The sub-group of users, which can be considered for uplink transmissions can, for
10 example, be decided based on the signal strength received by the macro users from the URLLC network.

In this way, the disclosures herein avoid scheduling cross-link (from macro downlink to URLLC uplink) interference as much as possible, i.e., when ever made possible by the load in the macro cells. The main motivation for this is that the eMBB traffic is typically
15 more flexible or less latency critical compared to the URLLC traffic, so it is possible, for example, to buffer the data, or avoid utilizing all possible downlink slots for eMBB transmissions. Note that the gains from the proposed approach increase in scenarios of low eMBB load.

Turning now to other embodiments, if the eMBB node (e.g. macro base station)
20 is applying beamforming, the above described embodiments may be enhanced by only avoiding transmitting some downlink beams during slots whereby the predicted URLLC TDD pattern is uplink. In general, only beams that are sufficiently strong within the coverage area of the URLLC network, e.g. exceeding a predefined threshold signal strength, may be classified as interfering beams, subject to scheduling restrictions. At the same time, beams
25 having a signal strength below such a threshold signal strength may be scheduled normally. In more detail, the set of interfering beams can for example consist of the beams pointing from the macro base station towards the local URLLC network, or beams that are serving macro users located close to the URLLC network, or beams which can measure the highest level of interference coming from the URLLC network. Alternatively, the URLLC network
30 and/or devices connected to the URLLC network can measure and detect beams transmitted by the overlapping macro cells, and inform the macro network about the detected beams and their signal strengths. Yet another alternative is to perform an initial analysis of interfering beams during the planning or tuning phases of the network deployment, and then to refine the classification of interfering beams during the operation phase of the networks.

35 In summary, where the node applies beamforming the step of scheduling 304 may comprise scheduling the eMBB TDD pattern of beams overlapping a coverage area of the URLLC with a strength above a threshold signal strength, based on the predicted URLLC

TDD pattern. E.g. in order to manage cross-link interference between the overlapping beams and the URLLC network.

Turning now to other embodiments, the embodiments above assume that (or provide optimal results when) the eMBB network and the URLLC network employ the same numerology and apply slot-based scheduling, e.g. if the slot length (and transmission time interval, TTI) is the same for both networks. However this isn't necessarily the case. If the two networks employ different numerologies or perform sub-slot scheduling then the slots or sub-slots may not necessarily align.

For example, consider a scenario where an eMBB network follows the 'normal' scheme with "Transmission time interval, TTI = slot = 14 symbols", but a URLLC factory network is following a "sub-slot scheduling" scheme, where the TTI is less than a slot, i.e., less than 14 symbols. For example, one slot of 14 symbols can contain 4 symbols for downlink (downlink sub-slot), 4 symbols for uplink (uplink sub-slot), and 6 guard symbols, (or similar). For this case, the networks are not generally "synchronized" (e.g., the TTIs do not align). The skilled person will appreciate that this is merely an example however and that other combinations of numerologies, slot and sub-slot scheduling may lead to unsynchronized slots/sub-slots.

Generally therefore, if the networks are using different numerologies, or if the URLLC network is applying sub-slot scheduling, the networks will have different TTI lengths.

In some embodiments, if the TTI lengths are different then the method 300 may further comprise adjusting one or other of the TTIs (e.g. the TTI of the URLLC network or the eMBB network) such that the TTIs are made equal, before the proposed mechanisms described above with respect to Fig. 3 are applied to manage cross-link interference. For example, the TTI length of the eMBB transmissions may be adjusted to match the TTI length of the URLLC network.

In more detail, in some embodiments the step of scheduling (304) may comprise adjusting a numerology associated with the TDD pattern of the eMBB network so as to synchronise at least one slot in the TDD pattern of the eMBB network with at least one slot in the TDD pattern of the URLLC network. Put another way, the eMBB network may change to the same numerology as the URLLC network.

In some embodiments, the step of scheduling (304) may comprise adjusting a sub-slot scheduling scheme of the eMBB network so as to synchronise at least one slot or sub-slot in the TDD pattern of the eMBB network with at least one slot or sub-slot in the TDD pattern of the URLLC network. As an example, the eMBB network may apply the same sub-slot scheduling scheme as the URLLC network.

In some embodiments, the step of scheduling (304) may comprise leaving at least one slot or sub-slot of the eMBB TDD pattern unutilised (e.g. blank) so as to manage

cross-link interference on a slot or sub-slot of the URLLC TDD pattern. Put another way, the eMBB may “blank” parts of the (downlink) slots. For example, the eMBB may blank the eMBB symbols that overlap with URLLC slots or sub-slots that may otherwise experience cross-link interference (e.g. that we want to protect).

5 In summary, embodiments herein propose a) a mechanism to predict URLLC TDD patterns (TDD pattern consisting of either slots or sub-slots depending on the URLLC scheduling scheme), b) a mechanism to manipulate (or apply scheduling restrictions for) eMBB TDD patterns (TDD pattern consisting of either slots or sub-slots) to avoid cross-link interference, c) a mechanism to “blank” some symbols of the eMBB slots to avoid cross-link
10 interference (TDD pattern consisting of slots, but with blanked symbols).

Turning now to Fig. 5 which shows a method 500 according to some embodiments herein. In this embodiment the node comprises a base station in a macro network (e.g. a macro base station) and the URLLC network comprises a factory network or similar. In a first block 502, the method comprises dividing the time horizon into windows T
15 comprising sub-windows $(t-k;t)$ and $(t;t+s)$ as described above with respect to Fig. 4. In a block 504 the node learns the historical URLLC activity and TDD pattern over the time window $(t-k;t)$. The method then comprises predicting 302; 506 a time division duplex, TDD, pattern for the URLLC based on the historical URLLC TDD data, over the time window $(t;t+s)$. Predicting a URLLC TDD pattern in this way was described above with respect to
20 block 302 of the method 300 and the detail therein will be understood to apply equally to the block 506.

In block 508 the method comprises the macro base station predicting its future traffic load, e.g. the traffic load of its associated UEs over the future time window $(t;t+s)$. Predicting the eMBB traffic load was described above with respect to the method 300 and
25 the detail therein will be understood to apply equally to block 508 of the method 500.

In block 510 the method comprises adjusting or scheduling the TDD pattern of one or more neighbouring base stations over a time interval $(t;t+s)$ in order to manage the cross-link interference experienced by the URLLC factory network. As described above, the scheduling is performed based on the predicted URLLC TDD pattern and the predicted
30 eMBB load. In this way cross-link interference may be managed and reduced to levels that ensure the latency and reliability requirements of the URLLC network may be met.

Turning now to other embodiments, as illustrated in Fig. 6, in some embodiments there is a node 600 in a communications network according to some embodiments herein. The node 600 is configured (e.g. adapted or programmed) to perform
35 any of the embodiments of the method 300 and/or the method 500 as described above.

Generally, the node 600 may comprise any component or network function (e.g. any hardware or software module) in the communications network suitable for performing the

functions described herein. For example, a node may comprise equipment capable, configured, arranged and/or operable to communicate directly or indirectly with a UE (such as a wireless device) and/or with other network nodes or equipment in the communications network to enable and/or provide wireless or wired access to the UE and/or to perform other functions (e.g., administration) in the communications network. Examples of nodes include, but are not limited to, access points (APs) (e.g., radio access points), base stations (BSs) (e.g., radio base stations, Node Bs, evolved Node Bs (eNBs) and NR NodeBs (gNBs)). Further examples of nodes include but are not limited to core network functions such as, for example, core network functions in a Fifth Generation Core network (5GC).

The node 600 may be configured or operative to perform the methods and functions described herein, such as embodiments of the method 300 or the method 500 as described above. The node 600 may comprise processing circuitry (or logic) 602. It will be appreciated that the node 600 may comprise one or more virtual machines running different software and/or processes. The node 600 may therefore comprise one or more servers, switches and/or storage devices and/or may comprise cloud computing infrastructure or infrastructure configured to perform in a distributed manner, that runs the software and/or processes.

The processor 602 may control the operation of the node 600 in the manner described herein. The processor 602 can comprise one or more processors, processing units, multi-core processors or modules that are configured or programmed to control the node 600 in the manner described herein. In particular implementations, the processor 602 can comprise a plurality of software and/or hardware modules that are each configured to perform, or are for performing, individual or multiple steps of the functionality of the node 600 as described herein.

The node 600 may comprise a memory 604. In some embodiments, the memory 604 of the node 600 can be configured to store program code or instructions that can be executed by the processor 602 of the node 600 to perform the functionality described herein. Alternatively or in addition, the memory 604 of the node 600, can be configured to store any requests, resources, information, data, signals, or similar that are described herein. The processor 602 of the node 600 may be configured to control the memory 604 of the node 600 to store any requests, resources, information, data, signals, or similar that are described herein.

It will be appreciated that the node 600 may comprise other components in addition or alternatively to those indicated in Fig. 6. For example, in some embodiments, the node 600 may comprise a communications interface. The communications interface may be for use in communicating with other nodes in the communications network, (e.g. such as other physical or virtual nodes). For example, the communications interface may be

configured to transmit to and/or receive from other nodes or network functions requests, resources, information, data, signals, or similar. The processor 602 of node 600 may be configured to control such a communications interface to transmit to and/or receive from other nodes or network functions requests, resources, information, data, signals, or similar.

5 Generally, the node 600 is for (e.g. configured or suitable for) scheduling enhanced mobile-broadband, eMBB, transmissions. In scenarios where the coverage area of the eMBB cell or beam interferes with an ultra-reliable low latency communication network, URLLC, the node may be configured to schedule eMBB transmissions in such a manner as to manage cross-link interference on the URLLC network.

10 Briefly, in one embodiment, the node 600 may be configured to predict a time division duplex, TDD, pattern for the URLLC based on historical URLLC TDD data. The node 600 may further be configured to schedule an eMBB TDD pattern, based on the predicted URLLC TDD pattern, so as to manage cross-link interference on the URLLC network.

15 Predicting a time division duplex, TDD, pattern for the URLLC based on historical URLLC TDD data was described in detail with respect to the method 300 in block 302 and the detail therein will be appreciated to apply equally to embodiments of the node 600.

20 Scheduling an eMBB TDD pattern, based on the predicted URLLC TDD pattern, so as to manage cross-link interference on the URLLC network was also described in detail with respect to the method 300 in block 304 and the detail therein will be appreciated to apply equally to embodiments of the node 600.

25 In some embodiments, for example, the processor is further caused to: predict traffic load for the eMBB based on historical eMBB traffic load data; and schedule the eMBB TDD pattern based on the predicted URLLC TDD pattern and the predicted eMBB traffic load.

In this way a node is provided that is capable of managing the cross-link interference on a URLLC network such that the URLLC network may meet its reliability and QoS requirements and thus facilitate emerging technologies that require low latency and extremely high reliability communications.

30 In another embodiment, there is provided a computer program product comprising a computer readable medium, the computer readable medium having computer readable code embodied therein, the computer readable code being configured such that, on execution by a suitable computer or processor, the computer or processor is caused to perform any of the embodiments of methods described herein, such as embodiments of the method 35 300 and/or the method 500.

Thus, it will be appreciated that the disclosure also applies to computer programs, particularly computer programs on or in a carrier, adapted to put embodiments into practice.

The program may be in the form of a source code, an object code, a code intermediate source and an object code such as in a partially compiled form, or in any other form suitable for use in the implementation of the method according to the embodiments described herein.

5 It will also be appreciated that such a program may have many different architectural designs. For example, a program code implementing the functionality of the method or system may be sub-divided into one or more sub-routines. Many different ways of distributing the functionality among these sub-routines will be apparent to the skilled person. The sub-routines may be stored together in one executable file to form a self-contained program. Such an executable file may comprise computer-executable instructions, for example, processor instructions and/or interpreter instructions (e.g. Java interpreter instructions). Alternatively, one or more or all of the sub-routines may be stored in at least one external library file and linked with a main program either statically or dynamically, e.g. at run-time. The main program contains at least one call to at least one of the sub-routines. The sub-routines may also comprise function calls to each other.

15 The carrier of a computer program may be any entity or device capable of carrying the program. For example, the carrier may include a data storage, such as a ROM, for example, a CD ROM or a semiconductor ROM, or a magnetic recording medium, for example, a hard disk. Furthermore, the carrier may be a transmissible carrier such as an electric or optical signal, which may be conveyed via electric or optical cable or by radio or other means. When the program is embodied in such a signal, the carrier may be constituted by such a cable or other device or means. Alternatively, the carrier may be an integrated circuit in which the program is embedded, the integrated circuit being adapted to perform, or used in the performance of, the relevant method.

25 Variations to the disclosed embodiments can be understood and effected by those skilled in the art in practicing the claimed invention, from a study of the drawings, the disclosure and the appended claims. In the claims, the word "comprising" does not exclude other elements or steps, and the indefinite article "a" or "an" does not exclude a plurality. A single processor or other unit may fulfil the functions of several items recited in the claims. The mere fact that certain measures are recited in mutually different dependent claims does not indicate that a combination of these measures cannot be used to advantage. 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. Any reference signs in the claims should not be construed as
30 limiting the scope.
35

CLAIMS

1. A method in a node of a communications network for scheduling enhanced mobile-broadband, eMBB, transmissions, wherein the eMBB coverage area interferes with an ultra-reliable low latency communication network, URLLC, the method comprising:
- 5 predicting (302) a time division duplex, TDD, pattern for the URLLC based on historical URLLC TDD data; and
- 10 scheduling (304) an eMBB TDD pattern, based on the predicted URLLC TDD pattern, so as to manage cross-link interference on the URLLC network.
2. A method as in claim 1 further comprising:
- 15 predicting traffic load for the eMBB based on historical eMBB traffic load data; and
- wherein the step of scheduling comprises scheduling the eMBB TDD pattern based on the predicted URLLC TDD pattern and the predicted eMBB traffic load.
3. A method as in claim 1 or 2 wherein the step of predicting (302) a time division duplex, TDD, pattern for the URLLC based on historical URLLC TDD data comprises predicting, at a time, t , the time division duplex, TDD, pattern for the URLLC over a future time interval, $(t;t+s)$, wherein s comprises the duration of the future time interval; and
- 20 wherein the prediction is based on historical URLLC TDD data obtained in a past time interval $(t-k;t)$, wherein k comprises the duration of the past time interval.
- 25
4. A method as in claim 3 when dependent on claim 2 wherein the step of scheduling (304) an eMBB TDD pattern comprises:
- 30 at the time t , allocating future eMBB TDD patterns over the future time interval $(t;t+s)$, based on the predicted URLLC TDD pattern over the future time interval $(t;t+s)$ and the predicted eMBB traffic load over the future time interval $(t;t+s)$.
5. A method as in any one of claims 1 to 4 wherein the step of scheduling (304) comprises:
- 35 scheduling the eMBB TDD pattern so as to meet a service level requirement associated with the URLLC.

6. A method as in any one of claims 1 to 5 wherein the step of scheduling (304) comprises:
- minimising a number of slots in which the eMBB transmits in the downlink whilst the URLLC is predicted to transmit in the uplink; and/or
 - 5 minimising a number of slots in which the eMBB transmits in the uplink whilst the URLLC is predicted to transmit in the downlink.
7. A method as in any one of claims 1 to 6 wherein the step of scheduling (304) comprises:
- 10 scheduling an eMBB uplink slot during a predicted URLLC uplink slot; or
 - scheduling an eMBB downlink slot during a predicted URLLC downlink slot.
8. A method as in any one of claims 1 to 7 wherein the step of scheduling (304) comprises:
- 15 refraining from scheduling an eMBB downlink slot during a predicted URLLC uplink slot; or
 - scheduling an eMBB downlink slot during a predicted URLLC uplink slot only if required by the eMBB network to meet a service level requirement associated with the eMBB.
- 20
9. A method as in any one of claims 1 to 8 wherein the step of scheduling (304) comprises:
- 25 scheduling an eMBB uplink slot during a predicted URLLC downlink slot for a subset of devices connected to the eMBB network, wherein the subset of devices are located such that the crosslink interference caused by uplink from the subset of devices on the downlink of the URLLC is below a threshold interference level.
10. A method as in claim 9 wherein the threshold interference level is set based on the signal strength of the subset of devices as measured in the URLLC network.
- 30
11. A method as in any one of claims 1 to 10 wherein the step of scheduling (304) comprises:
- 35 refraining from scheduling an eMBB slot during a predicted URLLC uplink or downlink slot, based on:
 - an amount of downlink data in the buffer;
 - a downlink latency requirement of the eMBB network; and/or
 - an average bit rate requirements of the eMBB.

12. A method as in any one of claims 1 to 11 wherein the node applies beamforming; and wherein the step of scheduling (304) comprises:
5 scheduling the eMBB TDD pattern of beams overlapping a coverage area of the URLLC with a strength above a threshold signal strength, based on the predicted URLLC TDD pattern.
13. A method as in any one of claims 1 to 12 wherein the step of scheduling (304) comprises:
10 adjusting a numerology associated with the TDD pattern of the eMBB so as to synchronise at least one slot in the TDD pattern of the eMBB with at least one slot in the TDD pattern of the URLLC; or
adjusting a sub-slot scheduling scheme of the eMBB so as to synchronise at
15 least one slot or sub-slot in the TDD pattern of the eMBB with at least one slot or sub-slot in the TDD pattern of the URLLC.
14. A method as in any one of claims 1 to 13 wherein the step of scheduling (304) comprises:
20 leaving at least one slot or symbol of the eMBB TDD pattern unutilised so as to manage cross-link interference on a slot or sub-slot of the URLLC.
15. A method as in any one of claims 1 to 14 wherein the step of predicting (302) comprises:
25 training a first model to predict the URLLC TDD pattern using a machine learning process.
16. A method as in claim 15 wherein the first model comprises a recurrent neural network.
- 30 17. A method as in claim 15 or 16 wherein the first model takes as input one or more of:
a time of day;
a day of the week;
a previous URLLC pattern;
a URLLC load; and
35 an estimate of a future load and/or TDD pattern provided by the URLLC network.

18. A method as in claim 15, 16 or 17 wherein training the first model comprises training the model on training data comprising i) example inputs and ii) corresponding ground truth URLLC TDD patterns.
- 5 19. A method as in claim 18 further comprising determining a ground truth URLLC TDD pattern based on interference experienced by the eMBB network from the URLLC network.
- 10 20. A method as in claim 19 wherein the URLLC and the eMBB are operated by different network operators.
21. A method as in claim 18 further comprising receiving a signal from the URLLC network, the signal comprising an indication of an example input and/or an example corresponding ground truth URLLC TDD pattern.
- 15 22. A method as in claim 21 wherein the URLLC network and the eMBB network are operated by the same network operator.
- 20 23. A method as in claim 2 wherein the step of predicting (302) traffic load for the eMBB comprises training a second model to predict the traffic load for the eMBB using a machine learning process.
24. A method as in claim 23 wherein the second model comprises a recurrent neural network.
- 25 25. A method as in claim 23 or 24 wherein the second model takes as input one or more of:
- a time of day;
 - a day of the week;
 - 30 a historical traffic load of the node associated with the eMBB; and
 - a historical traffic load associated with another node neighbouring the node associated with the eMBB.
- 35 26. A node (600) in a communications network, wherein the node is configured for enhanced mobile-broadband, eMBB, transmission with a coverage area that interferes with an ultra-reliable low latency communication network, URLLC, wherein the node comprises:

a memory (604) comprising instruction data representing a set of instructions;
and

a processor (602) configured to communicate with the memory and to execute the set of instructions, wherein the set of instructions, when executed by the processor,

5 cause the processor to:

predict a time division duplex, TDD, pattern for the URLLC based on historical URLLC TDD data; and

schedule an eMBB TDD pattern, based on the predicted URLLC TDD pattern, so as to manage cross-link interference on the URLLC network.

10

27. A node as in claim 26 wherein the processor is further caused to:

predict traffic load for the eMBB based on historical eMBB traffic load data; and

wherein being caused to schedule comprises the processor being caused to

schedule the eMBB TDD pattern based on the predicted URLLC TDD pattern and the

15

predicted eMBB traffic load.

28. A node as in claim 26 or 27 wherein the node comprises one of:

a base station;

an evolved NodeB; or

20

a network function in a 5G network.

29. A node as in any one of claims 26 to 28 wherein the node is configured to perform the method of any one of claims 3 to 25.

25

30. A computer program product comprising a computer readable medium, the computer readable medium having computer readable code embodied therein, the computer readable code being configured such that, on execution by a suitable computer or processor, the computer or processor is caused to perform the method as claimed in any one of claims 1 to 25.

30

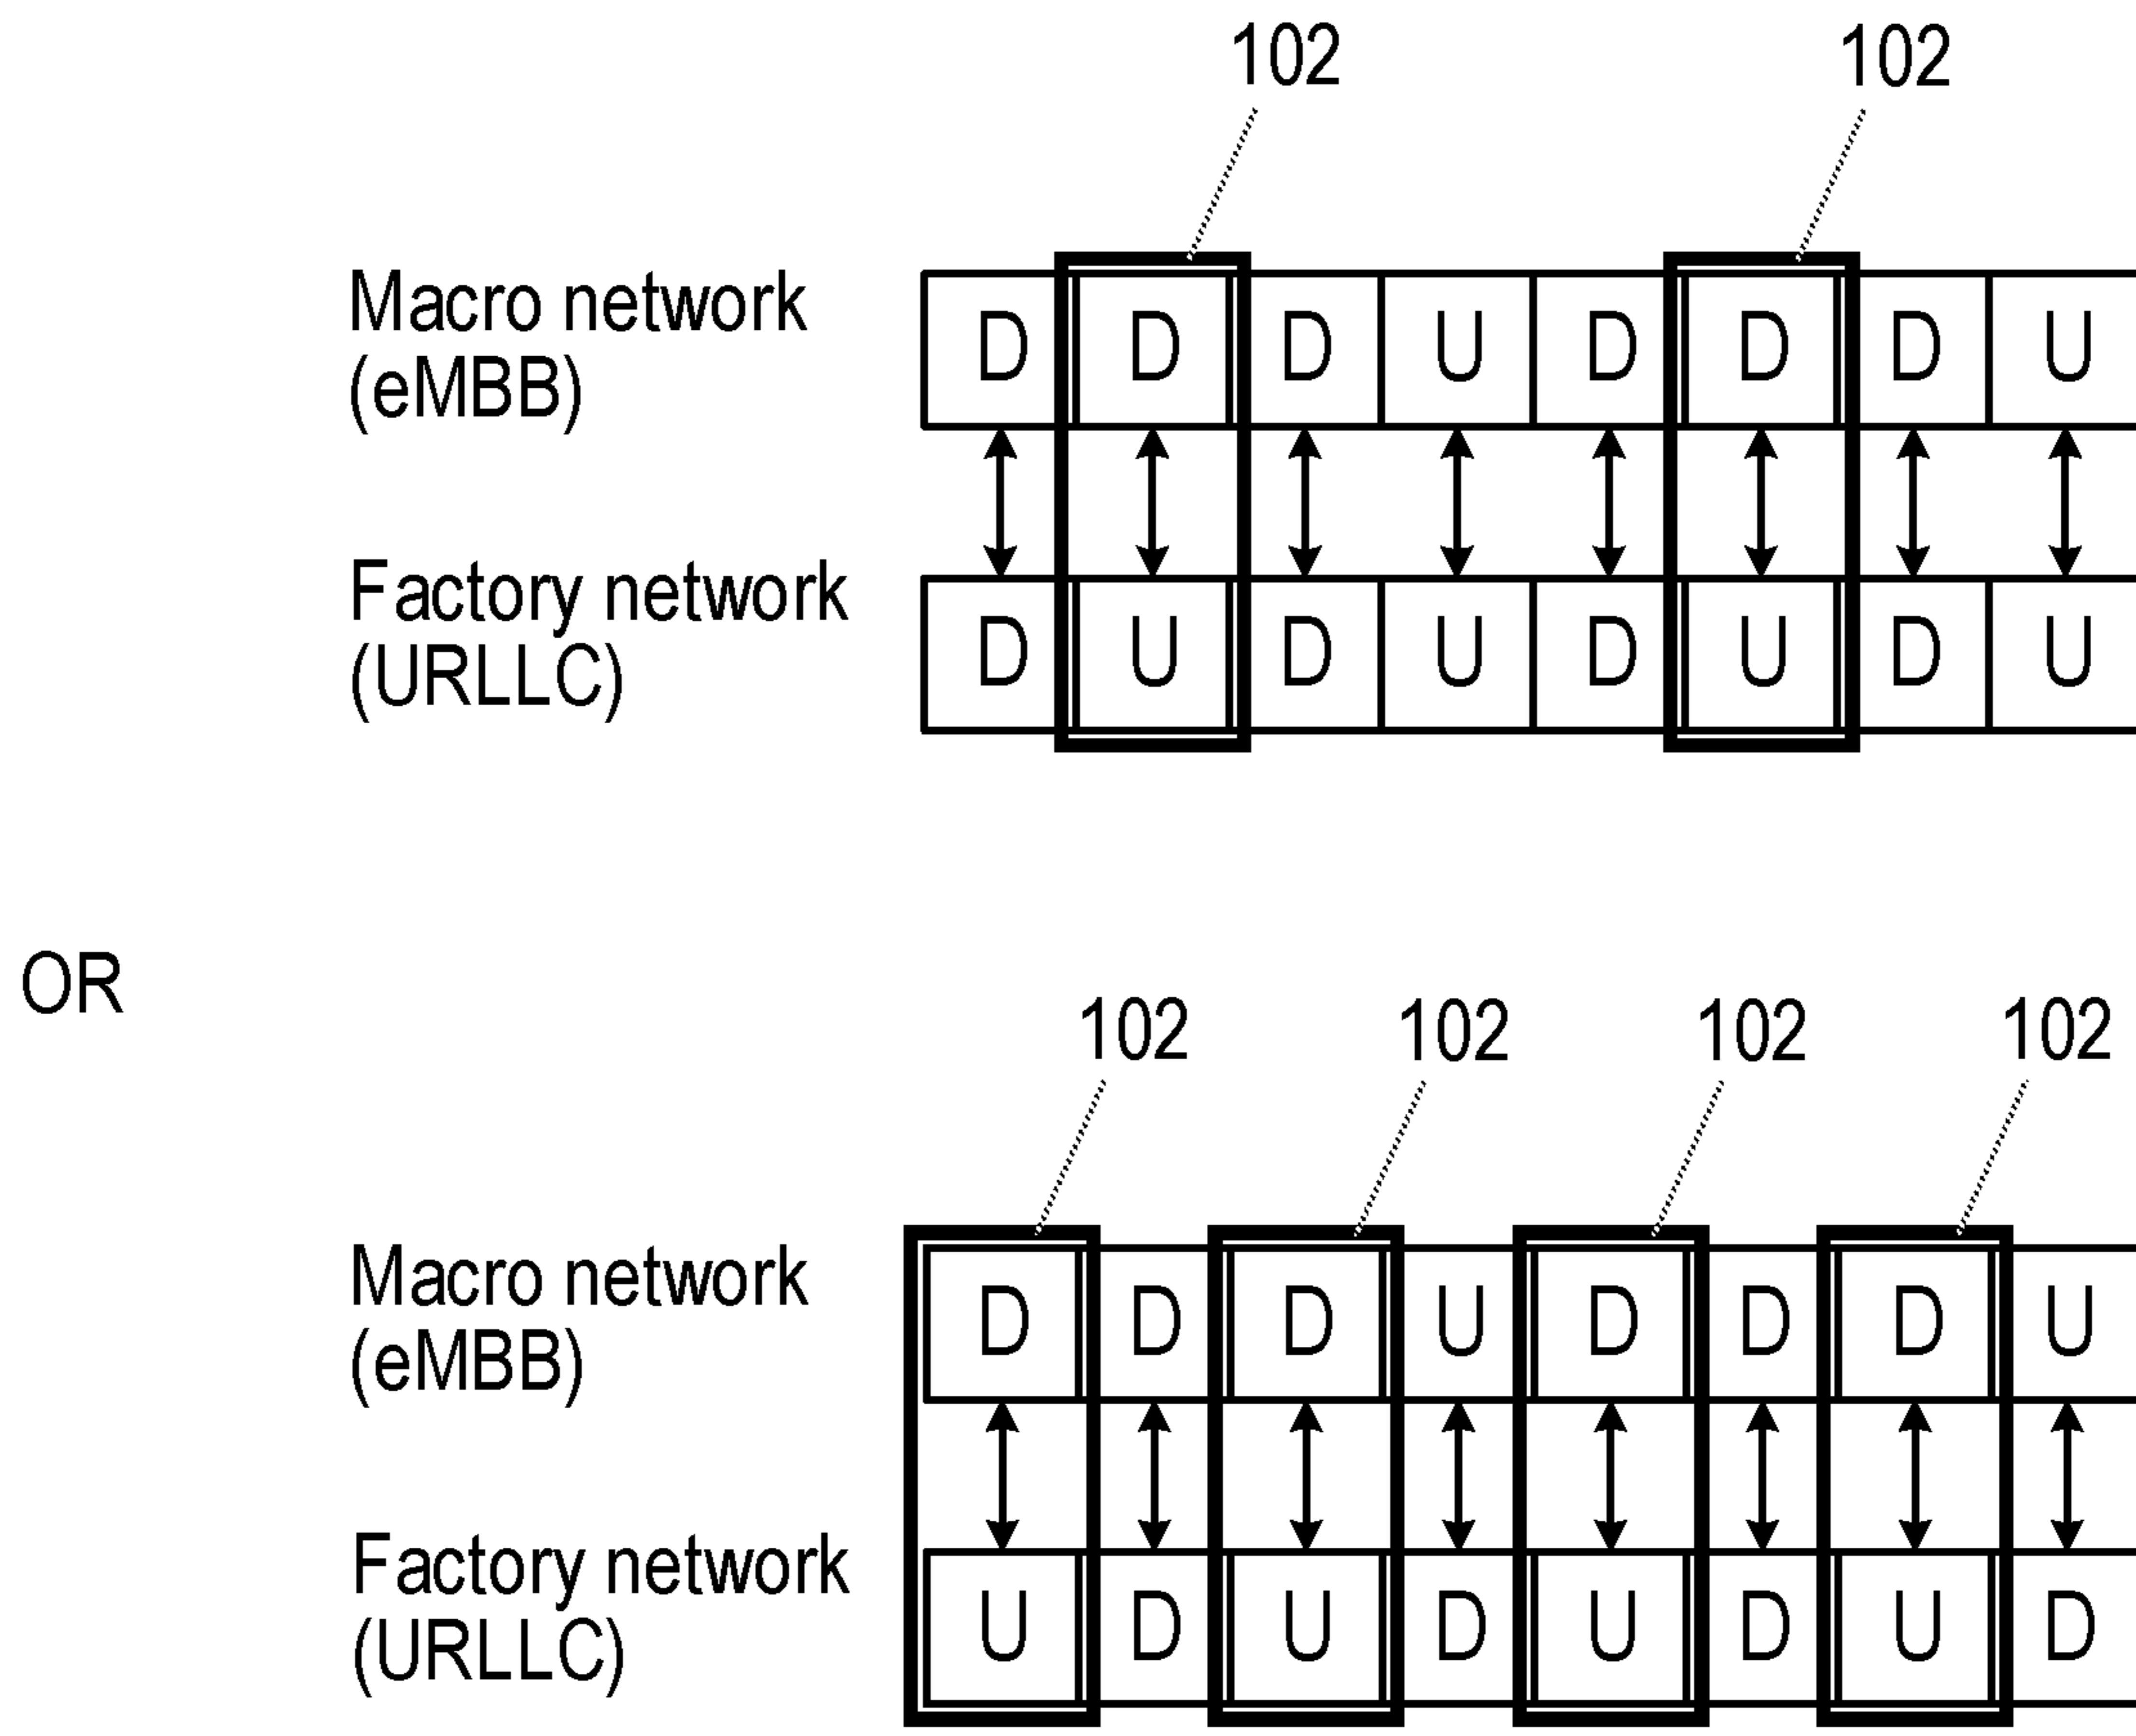


Fig. 1 (Prior Art)

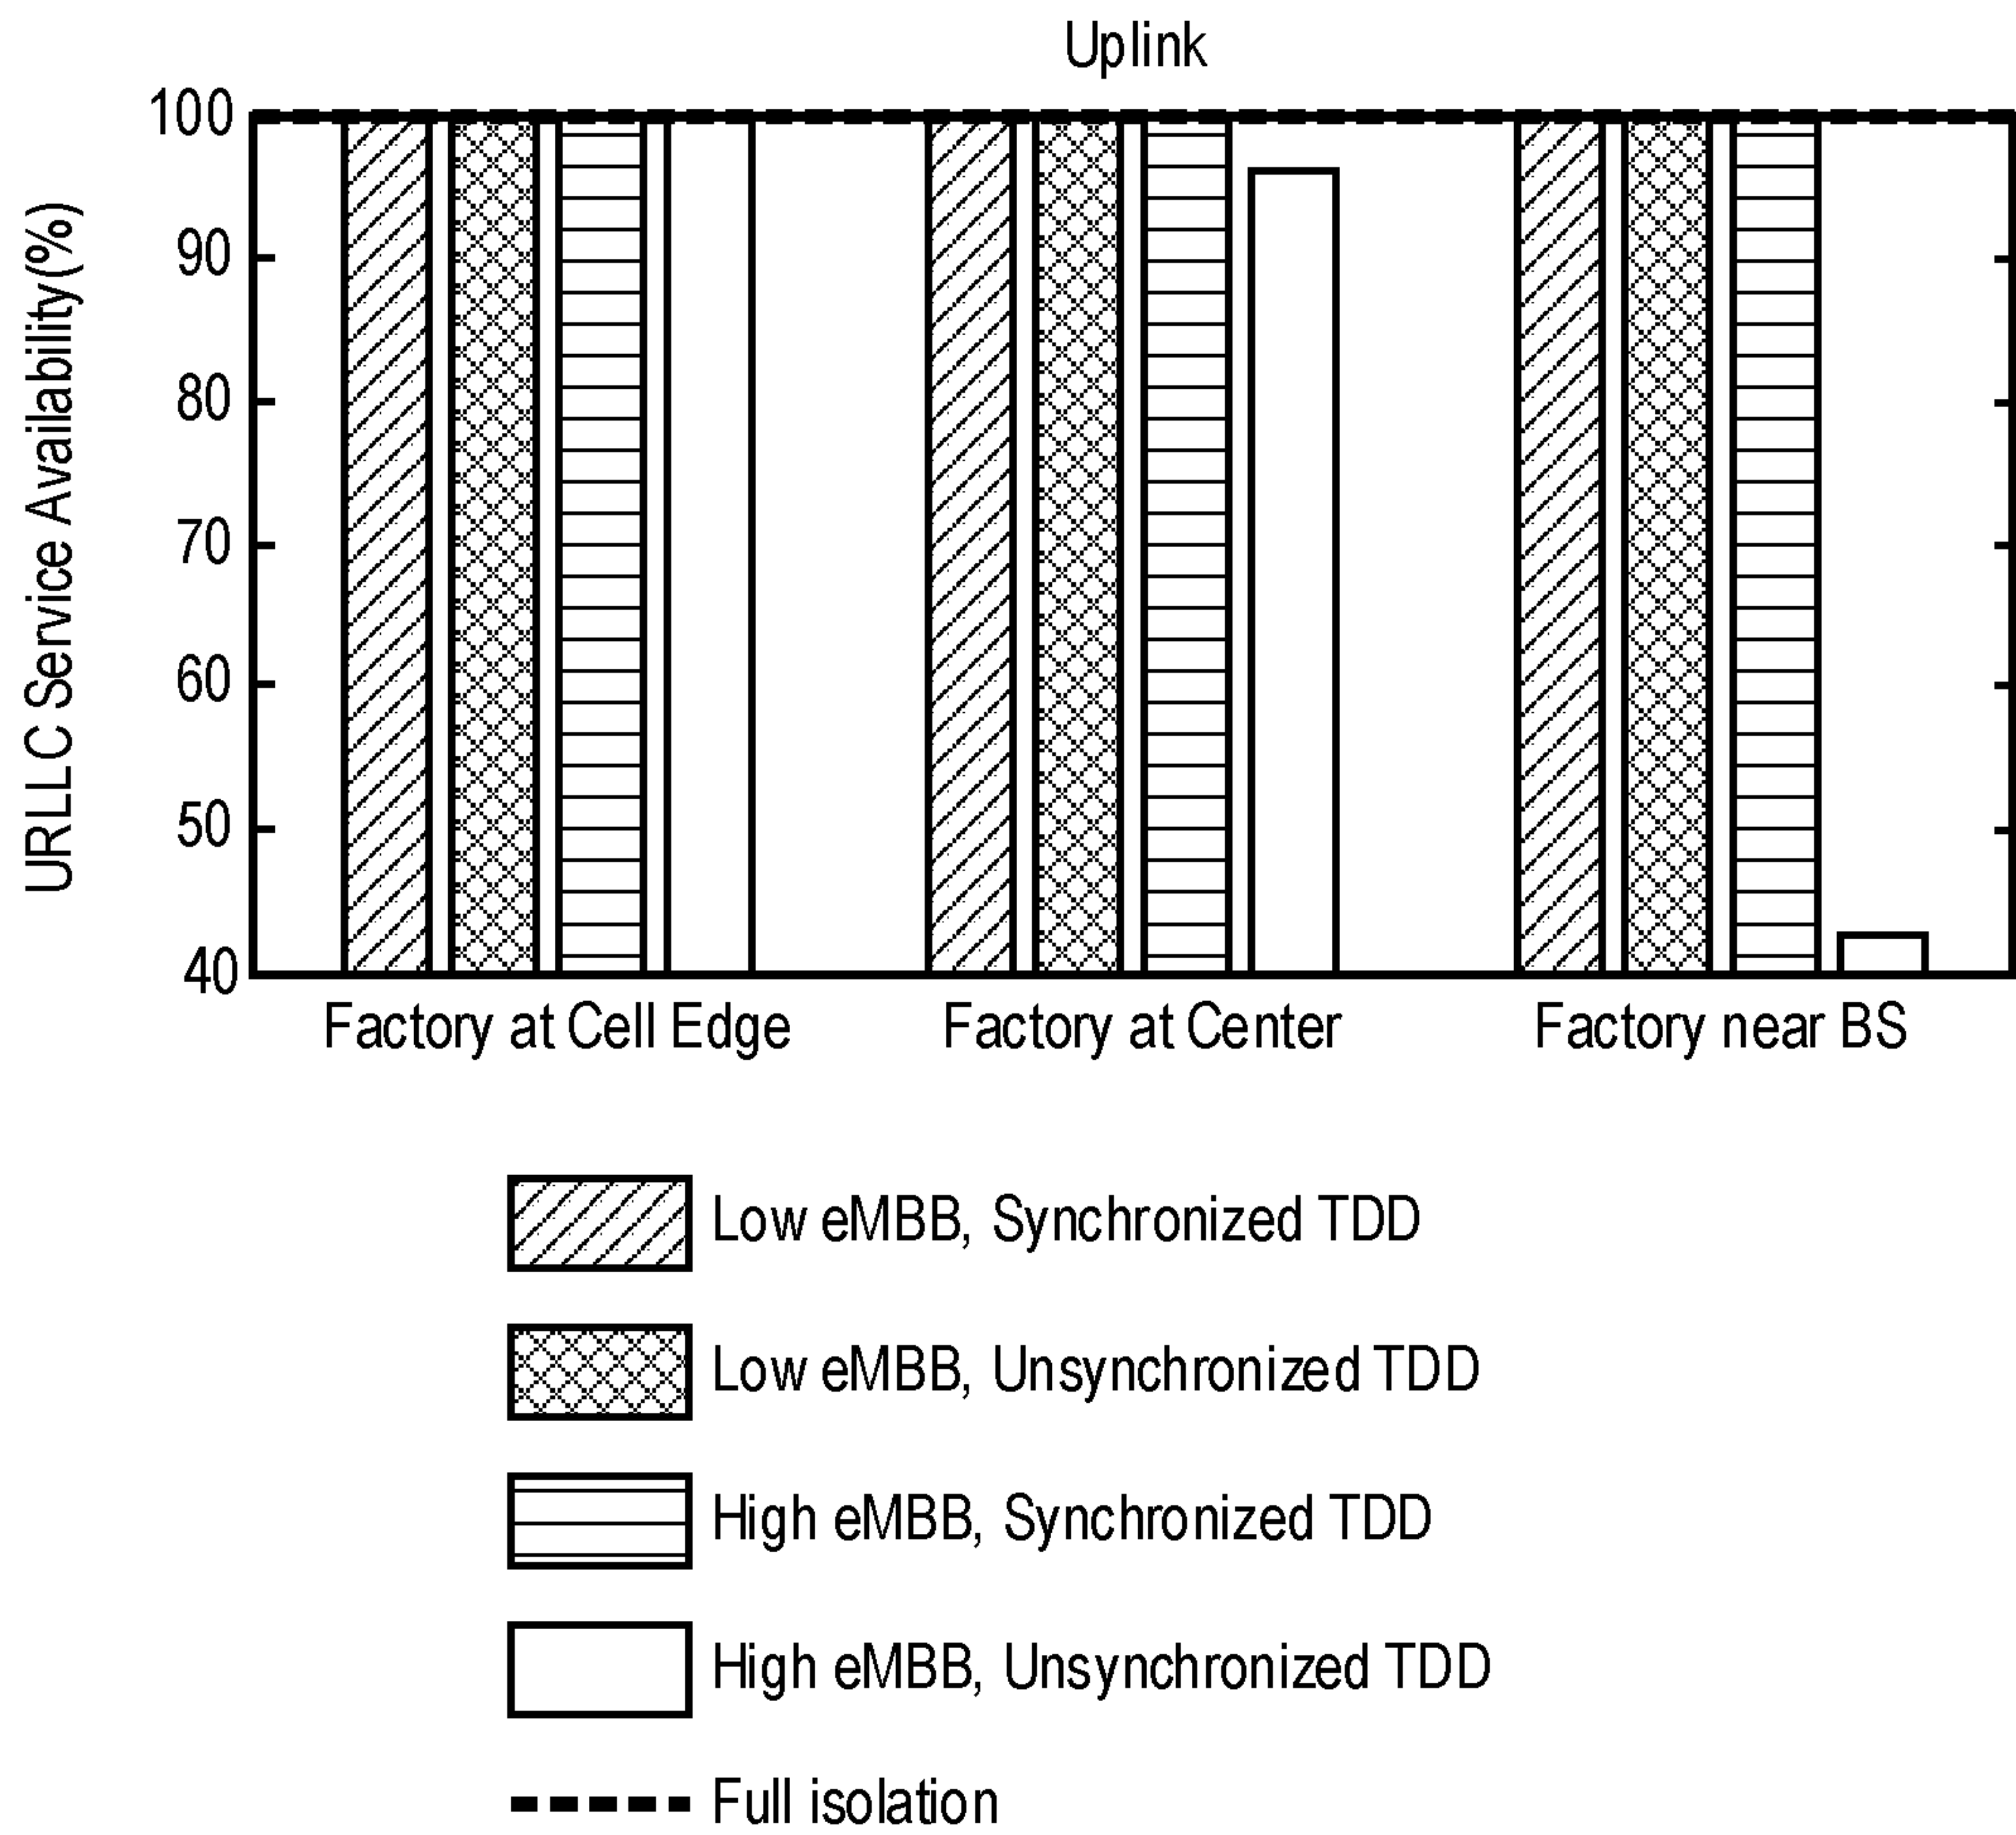


Fig. 2 (Prior Art)

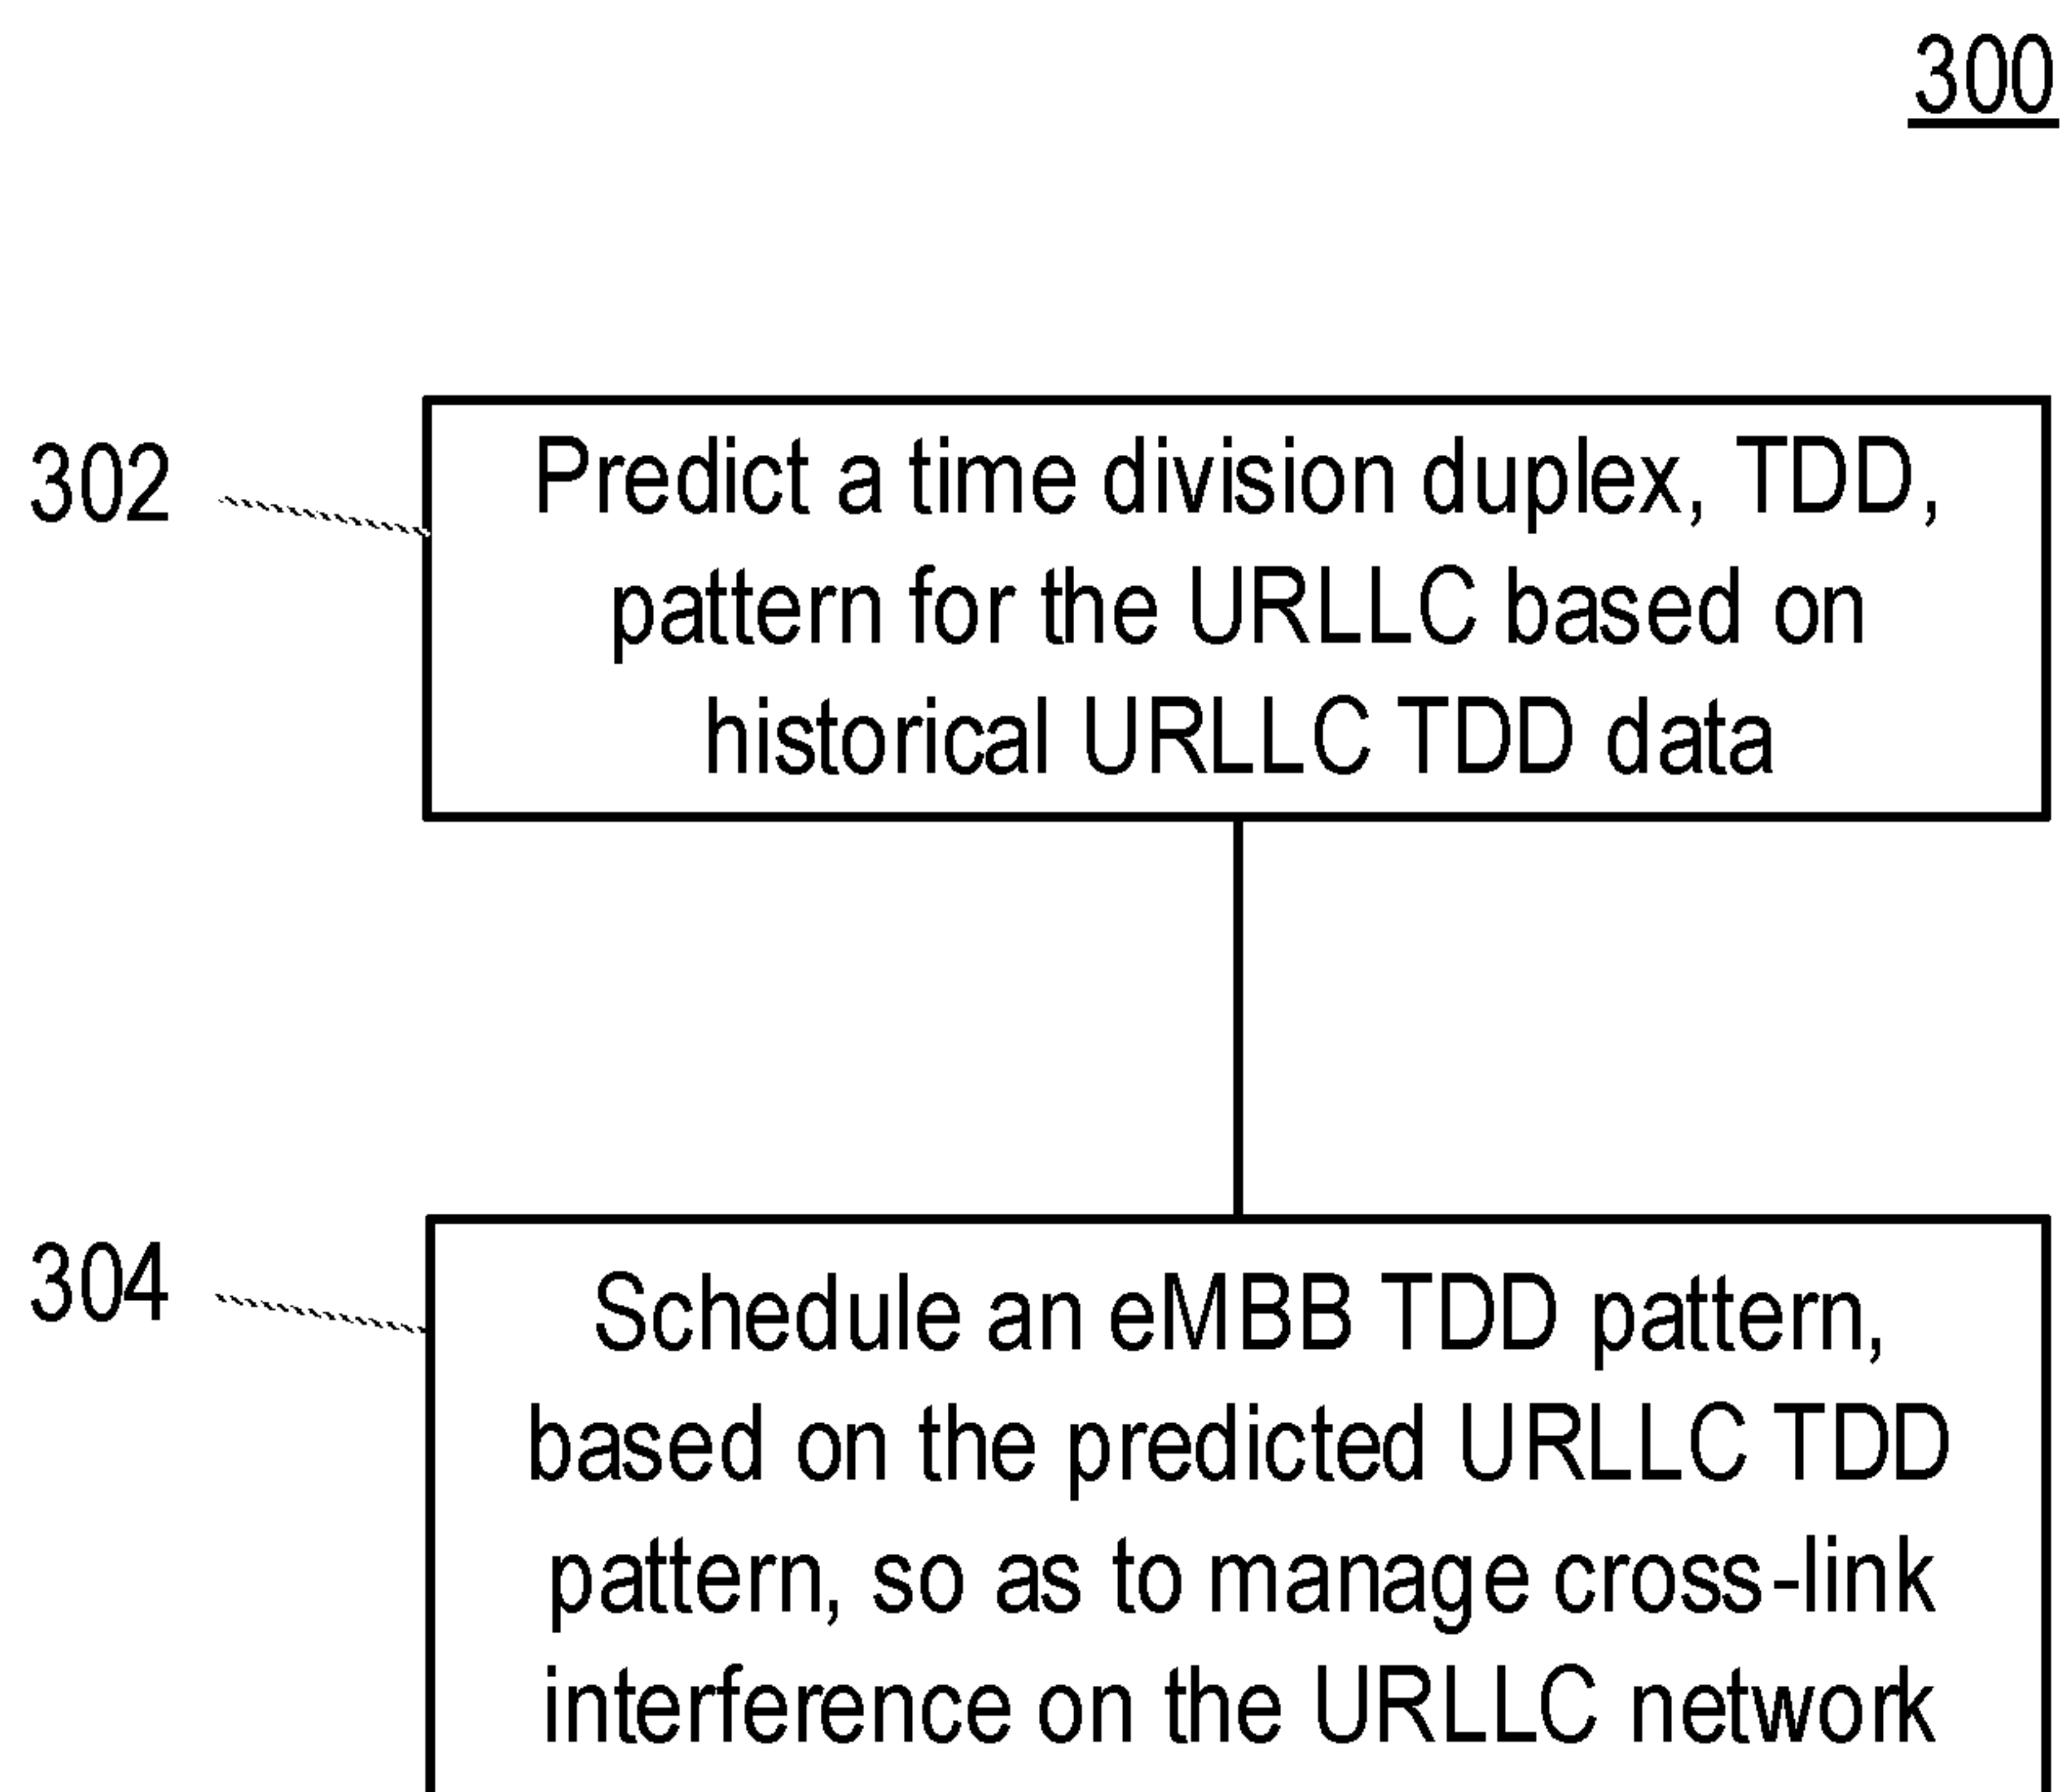


Fig. 3




Fig. 4

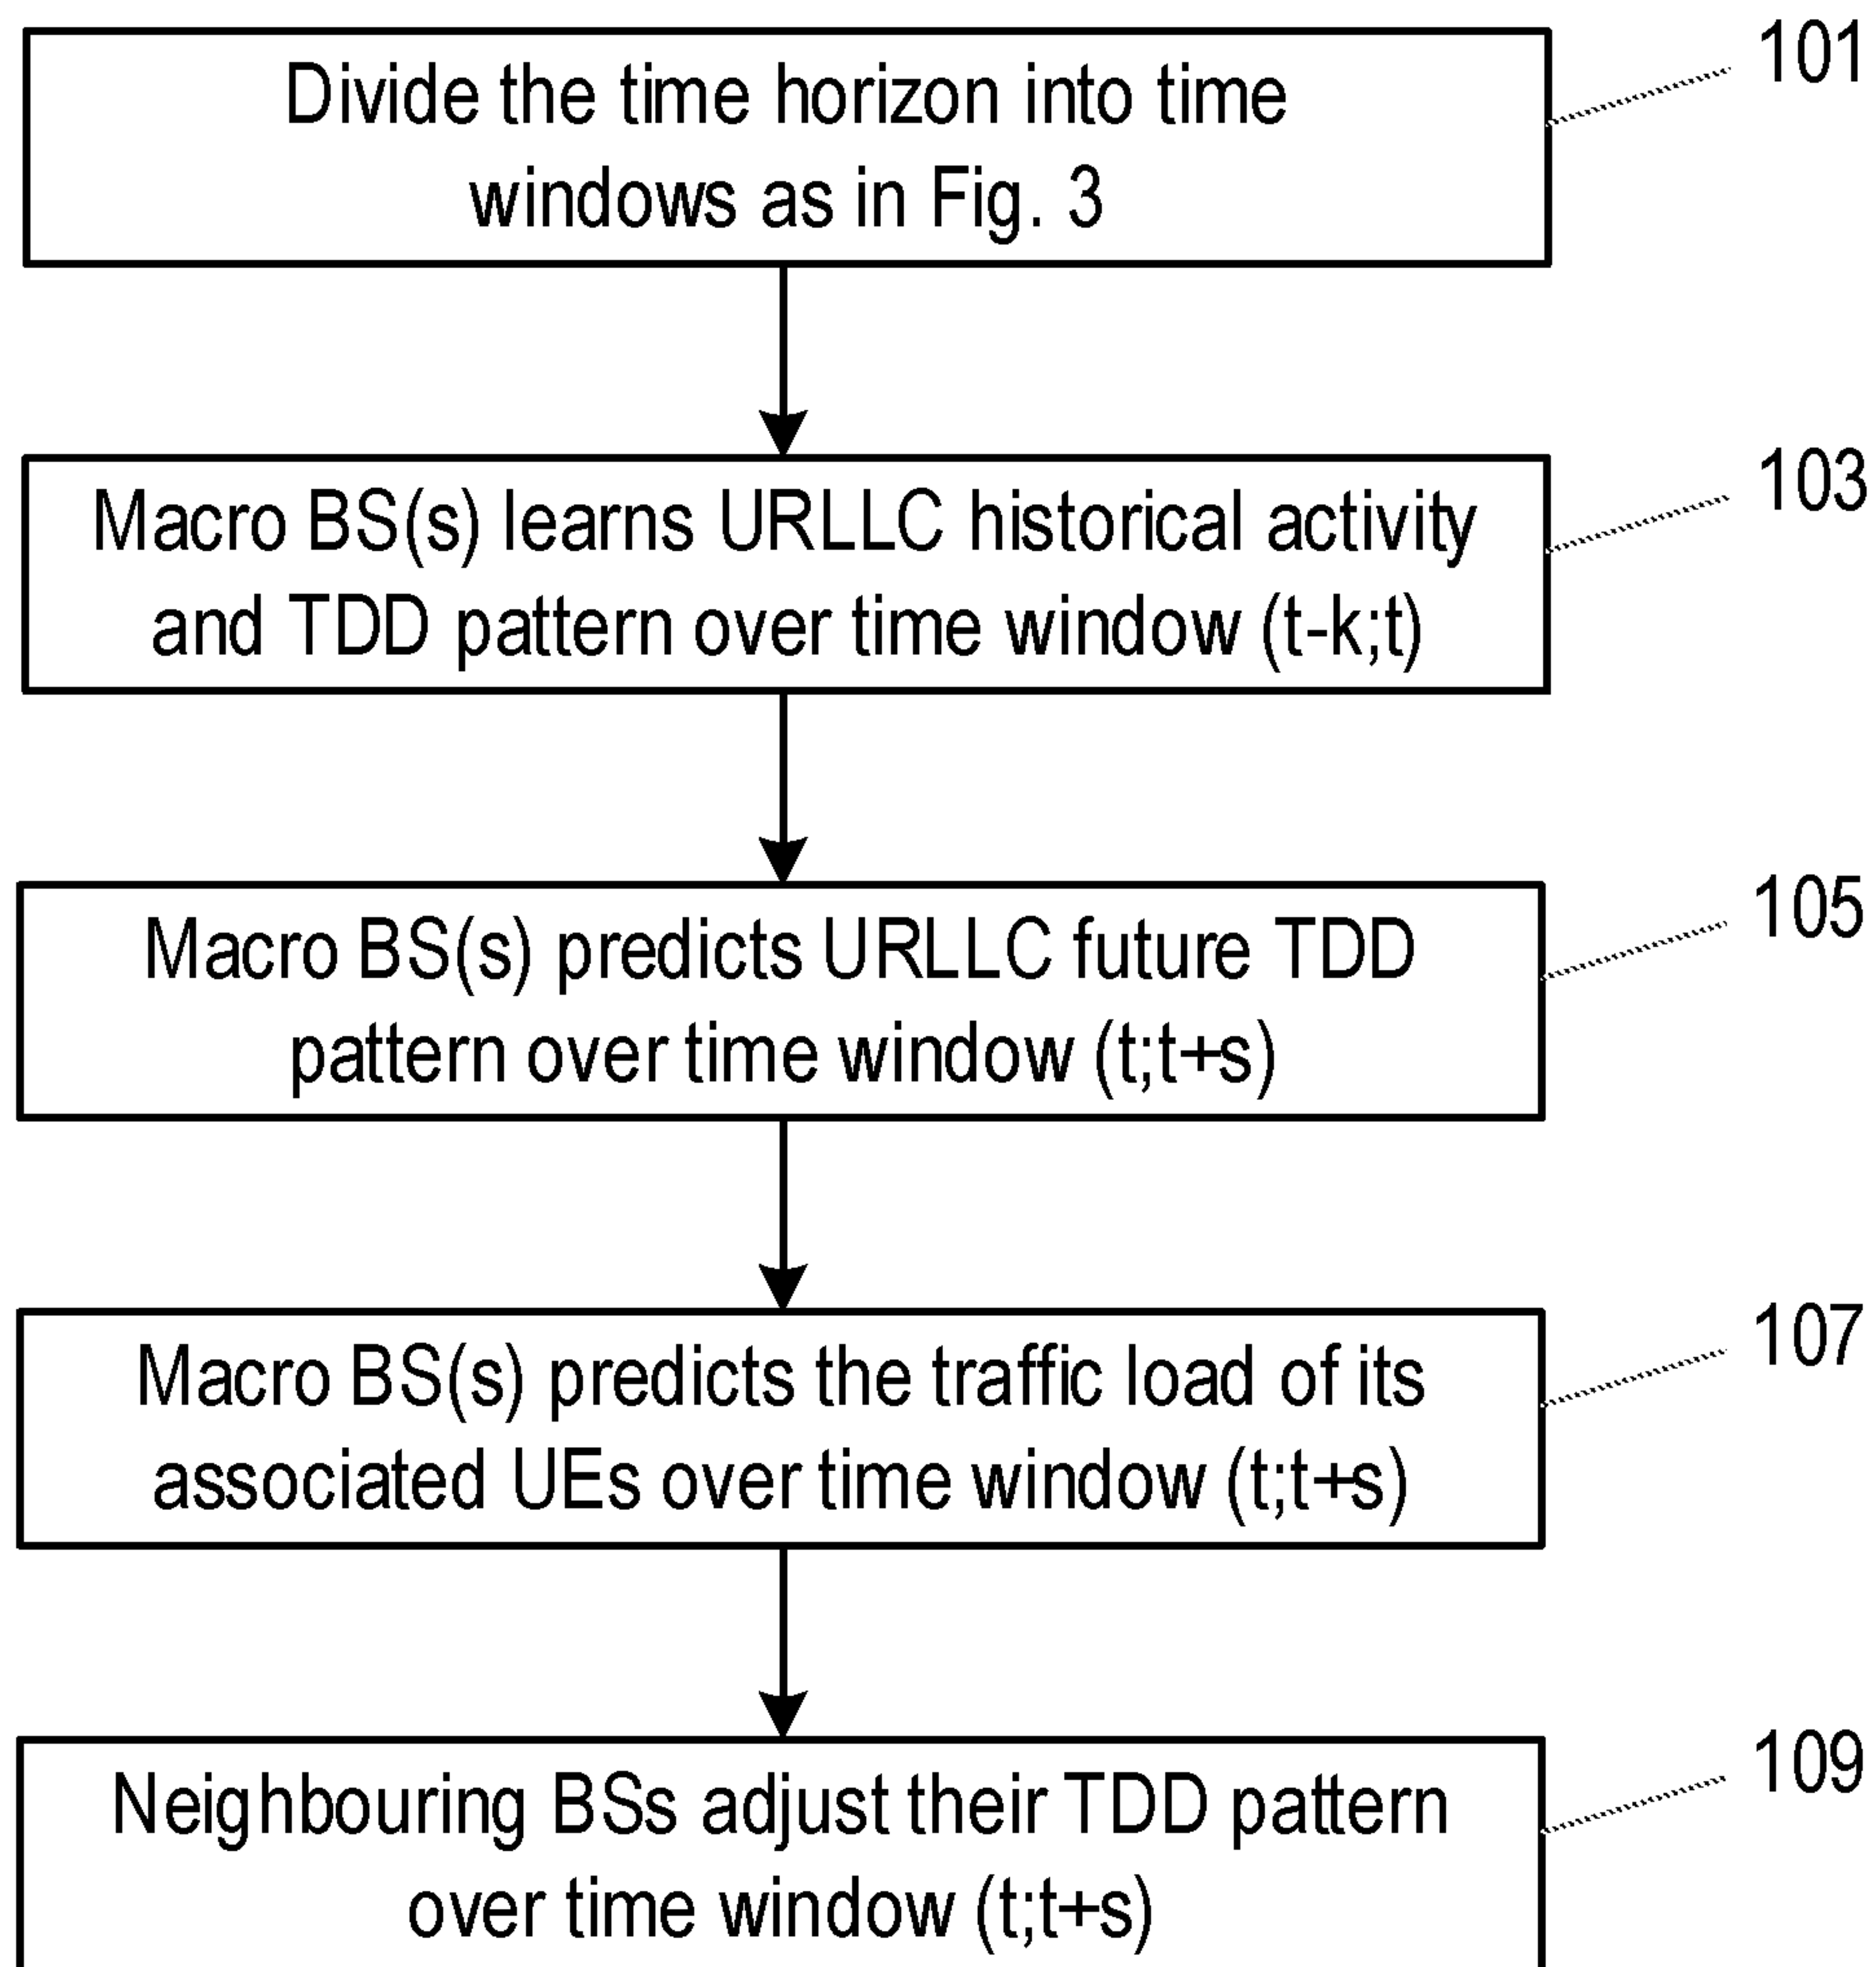


Fig. 5

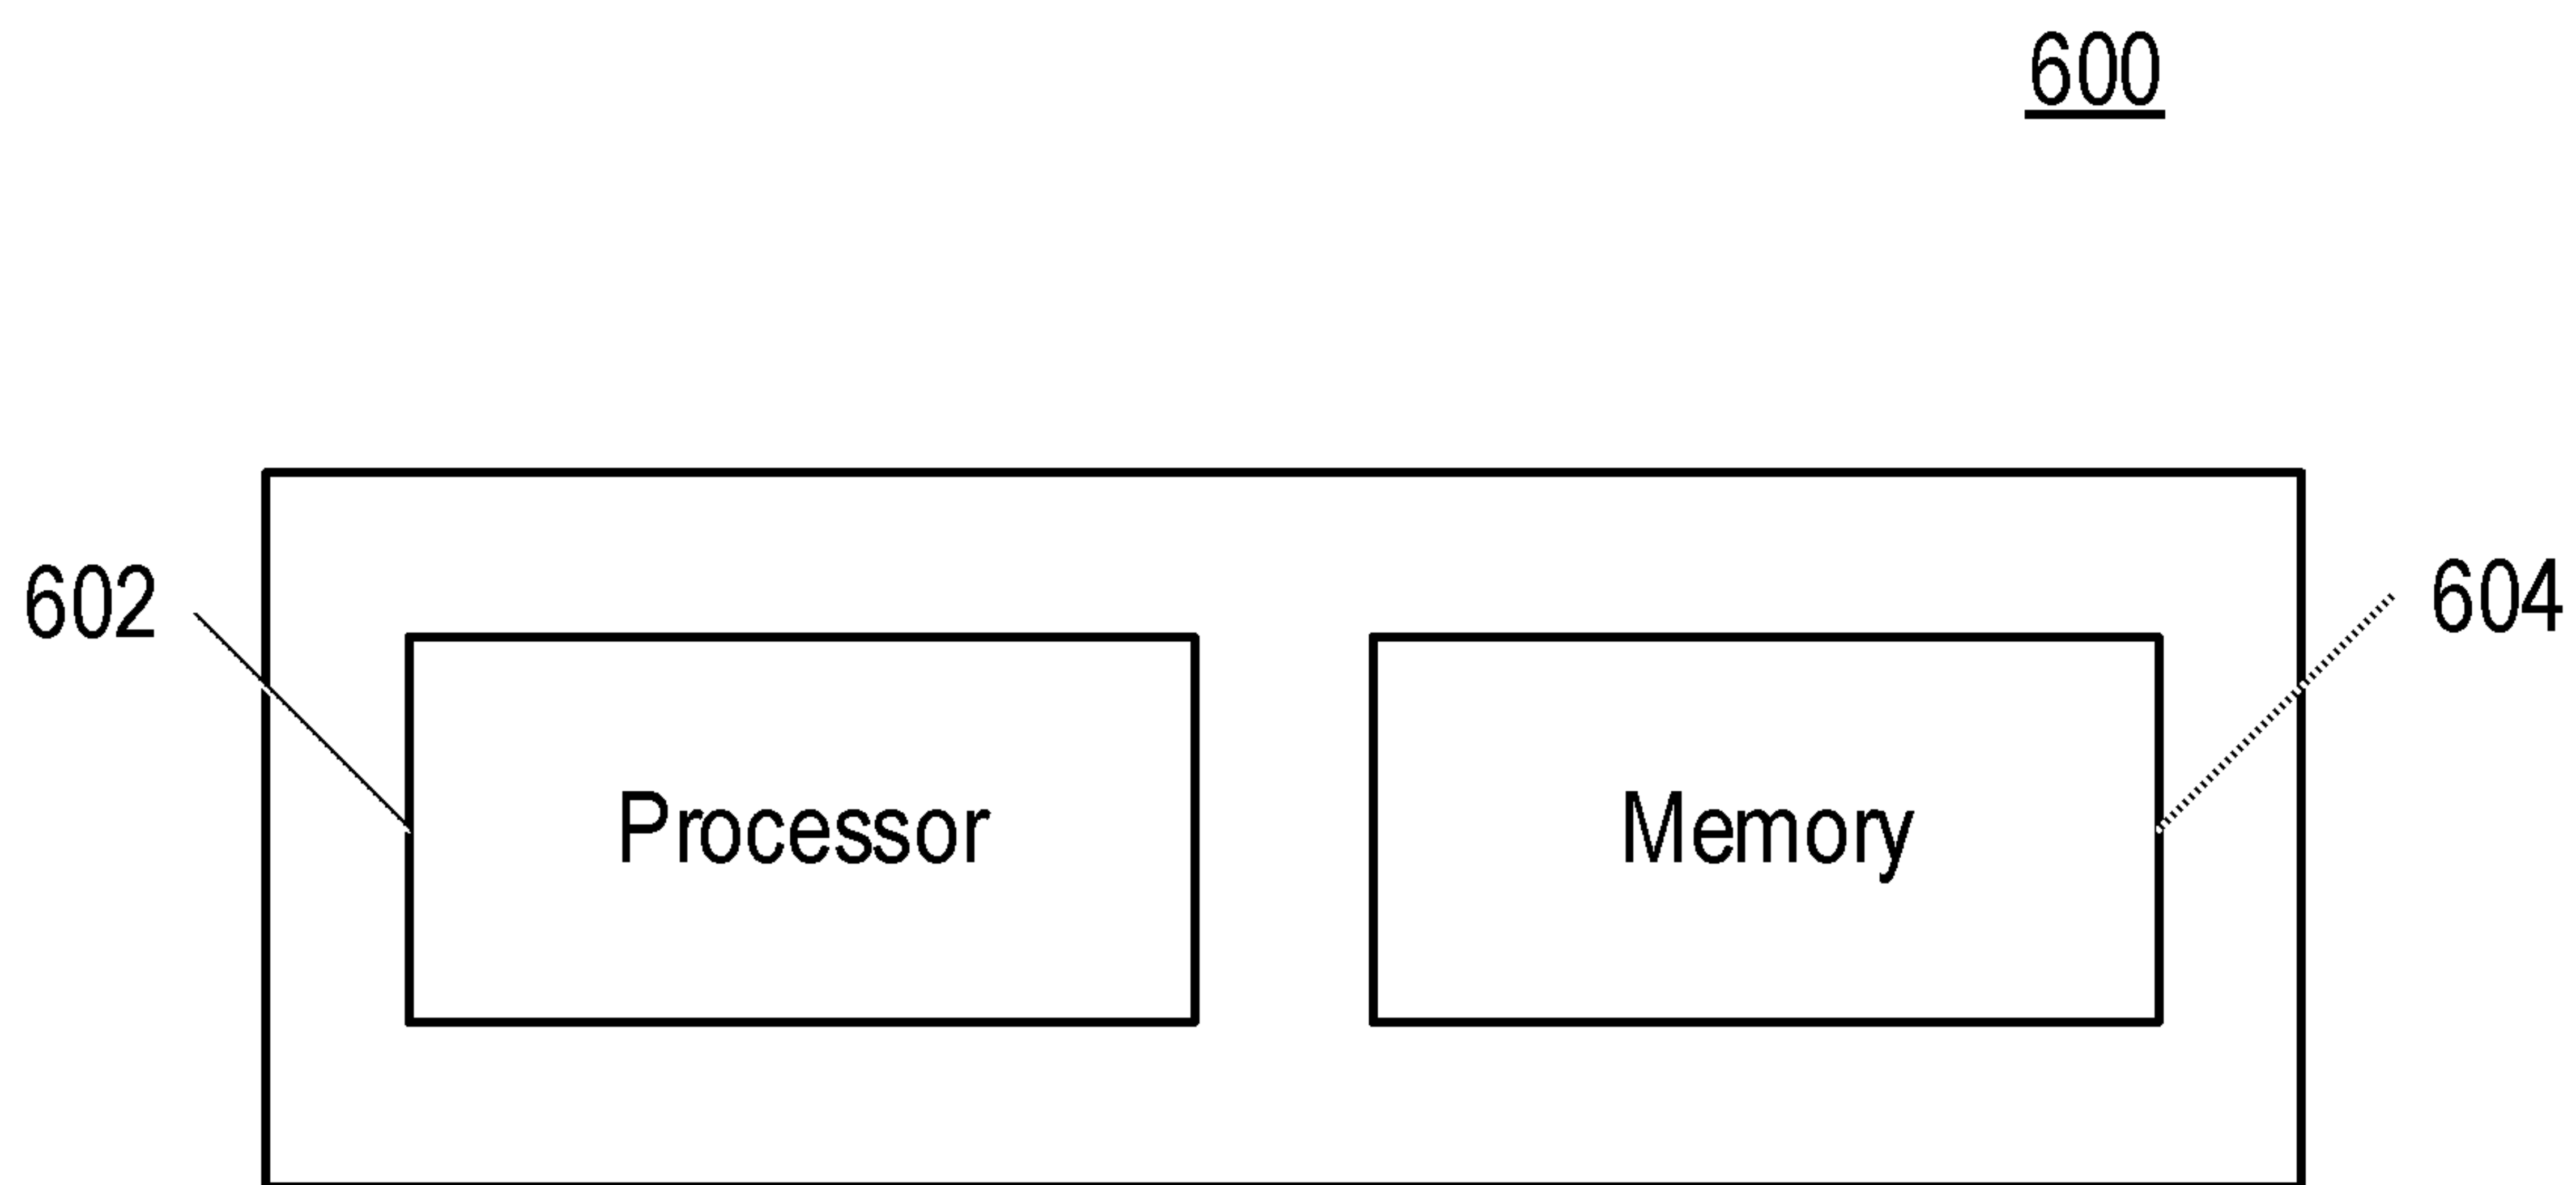


Fig. 6

INTERNATIONAL SEARCH REPORT

International application No.
PCT/SE2019/051124

A. CLASSIFICATION OF SUBJECT MATTER		
IPC: see extra sheet		
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)		
IPC: G06N, H04W		
Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched		
SE, DK, FI, NO classes as above		
Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)		
EPO-Internal, PAJ, WPI data, COMPENDEX, EMBASE, INSPEC		
C. DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	3GPP Draft; R1-1804450 Uplink Pre-emption for URLLC Reliability; 20180415; 3rd Generation Partnership Project (3GPP), Mobile Competence Centre ; 650, route des Lucioles ; F-06921 Sophia-Antipolis Cedex ; France; Sanya, China; 20180416 - 20180420 http://www.3gpp.org/ftp/Meetings%5F3GPP%5FSYNC/RAN1/Docs/ ; whole document --	1-30
<input checked="" type="checkbox"/> Further documents are listed in the continuation of Box C. <input checked="" type="checkbox"/> 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 "D" document cited by the applicant in the international application "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 19-10-2020		Date of mailing of the international search report 20-10-2020
Name and mailing address of the ISA/SE Patent- och registreringsverket Box 5055 S-102 42 STOCKHOLM Facsimile No. + 46 8 666 02 86		Authorized officer Per Karlsson Telephone No. + 46 8 782 28 00

INTERNATIONAL SEARCH REPORT

International application No.
PCT/SE2019/051124

C (Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	A Matching Based Coexistence Mechanism between eMBB and uRLLC in 5G Wireless Networks; Anupam Kumar Bairagi, Shirajum Munir, Madyan Alsenwi, Nguyen H. Tran, Choongseon Hong; SAC '19: Proceedings of the 34th ACM/SIGAPP Symposium on Applied Computing April 2019 Pages 2377-2384; https://doi.org/10.1145/3297280.3297513 ; whole document; abstract; section 3 --	1-30
A	WO 2018129325 A1 (IDAC HOLDINGS INC), 12 July 2018 (2018-07-12); abstract; paragraphs [0119]-[0150]; figures 11-18 --	1-30
A	3GPP Draft; R1-1716036; Cross-link interference management based on coordinated beamforming; 20170917; 3rd Generation Partnership Project (3GPP), Mobile Competence Centre ; 650, route des Lucioles ; F-06921 Sophia-Antipolis Cedex ; France; Nagoya, Japan; 20170918 - 20170921 http://www.3gpp.org/ftp/Meetings_3GPP_SYNC/RAN1/Docs/ ; whole document --	1-30
A	3GPP Draft; R1-1710781; UE behaviour related to dynamic and semi-static configured Resources; 20170617; Qingdao, P.R. China; 20170627 - 20170630 http://www.3gpp.org/ftp/tsg_ran/WG1_RL1/TSGR1_AH/NR_AH_1706/Docs/ ; whole document --	1-30
A	Performance Evaluation for the Co-existence of eMBB and URLLC Networks: Synchronized versus Unsynchronized TDD; Ursula Challita, Kimmo Hiltunen, Miurel Tercero; publ 2019-9-22; 2019 IEEE 90th Vehicular Technology Conference (VTC2019-Fall) https://arxiv.org/ftp/arxiv/papers/1906/1906.00287.pdf ; whole document -- -----	1-30

Continuation of: second sheet

International Patent Classification (IPC)

H04W 72/10 (2009.01)

G06N 20/00 (2019.01)

H04W 72/14 (2009.01)

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/SE2019/051124

WO	2018129325 A1	12/07/2018	CN	110326349 A	11/10/2019
			EP	3566516 A1	13/11/2019
			US	20190327757 A1	24/10/2019
