

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

(19) World Intellectual Property

Organization

International Bureau

(43) International Publication Date

16 February 2023 (16.02.2023)



(10) International Publication Number

WO 2023/018933 A1

(51) International Patent Classification:

H04W 28/08 (2009.01) H04W 4/02 (2018.01)

H04W 40/24 (2009.01) H04L 47/125 (2022.01)

(21) International Application Number:

PCT/US2022/040157

(22) International Filing Date:

12 August 2022 (12.08.2022)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

17/400,943 12 August 2021 (12.08.2021) US

(71) Applicant: **DISH WIRELESS L.L.C.** [US/US]; 9601 S. Meridian Blvd., Englewood, Colorado 80112 (US).

(72) Inventors: **ALASTI, Mehdi**; c/o DISH Wireless L.L.C., 9601 S. Meridian Blvd., Englewood, Colorado 80112 (US). **BASHIR, Kazi**; DISH Wireless L.L.C., 9601 S. Meridian Blvd., Englewood, Colorado 80112 (US). **KHAMAS, Ash**; c/o DISH Wireless L.L.C., 9601 S. Meridian Blvd., Englewood, Colorado 80112 (US). **BANSAL, Ashish**; DISH Wireless L.L.C., 9601 S. Meridian Blvd., Englewood, Colorado 80112 (US). **CHENUMOLU, Siddhartha**; DISH Wireless L.L.C., 9601 S. Meridian Blvd., Englewood, Colorado 80112 (US).

(74) Agent: **BAUNACH, Jeremiah J.**; c/o VLP Law Group, LLP, 555 Bryant Street Suite 820, Palo Alto, California 94301 (US).

(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, CV, CZ, DE, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IQ, IR, IS, IT, JM, 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: USER PLANE FUNCTION (UPF) LOAD BALANCING BASED ON CURRENT UPF LOAD AND THRESHOLDS THAT DEPEND ON UPF CAPACITY

(57) Abstract: Embodiments are directed towards systems and methods for user plane function (UPF) and network slice load balancing within a 5G network. Example embodiments include systems and methods for load balancing based on current UPF load and thresholds that depend on UPF capacity; UPF load balancing using predicted throughput of new UE on the network based on network data analytics; UPF load balancing based on special considerations for low latency traffic; UPF load balancing supporting multiple slices, maintaining several load-thresholds for each UPF and each slice depending on the UPF and network slice capacity; and UPF load balancing using predicted central processing unit (CPU) utilization and/or predicted memory utilization of new UE on the network based on network data analytics.



WO 2023/018933 A1

USER PLANE FUNCTION (UPF) LOAD BALANCING BASED ON CURRENT UPF  
LOAD AND THRESHOLDS THAT DEPEND ON UPF CAPACITY

Technical Field

The present disclosure relates generally to digital message communications  
5 and, more particularly, to user plane function (UPF) load balancing within a Fifth  
Generation (5G) communications network.

BRIEF SUMMARY

As the use of smart phones and Internet of Things (IoT) devices has  
increased, so too has the desire for more reliable, fast, and continuous transmission  
10 of content. In an effort to improve the content transmission, networks continue to  
improve with faster speeds and increased bandwidth. The advent and  
implementation of 5G has resulted in faster speeds and increased bandwidth, but  
with the drawback of potentially overloading certain portions of the network in certain  
circumstances. It is with respect to these and other considerations that the  
15 embodiments described herein have been made.

5G Core (5GC) is the heart of a 5G mobile network. It establishes reliable,  
secure connectivity to the network for end users and provides access to its services.  
The core domain handles a wide variety of essential functions in the mobile network,  
such as connectivity of new user equipment (UE) and mobility management,  
20 authentication and authorization, subscriber data management and policy  
management, among others. 5G Core network functions are completely software-  
based and designed as cloud-native, meaning that they're agnostic to the underlying  
cloud infrastructure, allowing higher deployment agility and flexibility.

With the advent of 5G, industry experts defined how the core network should  
25 evolve to support the needs of 5G New Radio (NR) and the advanced use cases  
enabled by it. Together, they developed the 3rd Generation Partnership Project  
(3GPP) standard for core networks known as 5G Core (5GC).

The 5GC architecture is based on what is called a Service-Based Architecture  
(SBA), which implements IT network principles and a cloud-native design approach.  
30 In this architecture, each network function (NF) offers one or more services to other  
NFs via Application Programming Interfaces (API). Each NF, such as the user plane

function (UPF) and the Session Management Function (SMF) is formed by a combination of small pieces of software code called as microservices. Some microservices can even be re-used for different NFs, making implementation more effective and facilitating independent life-cycle management – which allows  
5 upgrades and new functionalities to be deployed with zero impact on running services.

Briefly described, embodiments are directed toward systems and methods for user plane function (UPF) and network slice load balancing within a 5G network. Example embodiments include: systems and methods for load balancing based on  
10 current UPF load and thresholds that depend on UPF capacity; systems and methods for UPF load balancing using predicted throughput of new UE on the network based on network data analytics; systems and methods for UPF load balancing based on special considerations for low latency traffic; systems and methods for UPF load balancing supporting multiple slices, maintaining several load-  
15 thresholds for each UPF and each slice depending on the UPF and network slice capacity; and systems and methods for UPF load balancing using predicted central processing unit (CPU) utilization and/or predicted memory utilization of new UE on the network based on network data analytics.

#### BRIEF DESCRIPTION OF THE SEVERAL VIEWS OF THE DRAWINGS

20 Non-limiting and non-exhaustive embodiments are described with reference to the following drawings. In the drawings, like reference numerals refer to like parts throughout the various figures unless otherwise specified.

For a better understanding of the present invention, reference will be made to the following Detailed Description, which is to be read in association with the  
25 accompanying drawings:

Figure 1 illustrates a context diagram of an environment in which UPF load balancing may be implemented in accordance with embodiments described herein;

Figure 2 illustrates a logical flow diagram showing one embodiment of a process for load balancing based on current UPF load and thresholds that depend  
30 on UPF capacity in accordance with embodiments described herein;

Figure 3 illustrates a logical flow diagram showing one embodiment of a process for selecting the UPF based on generated weights, which is useful in the process of Figure 2 in accordance with embodiments described herein;

Figure 4 illustrates a logical flow diagram showing one embodiment of a process for selecting the UPF based on the determined load-regions for a plurality of UPFs and the weights generated based on the determined load regions, which is useful in the process of Figure 3 in accordance with embodiments described herein;

5           Figure 5 illustrates a logical flow diagram showing one embodiment of a process for UPF load balancing using predicted throughput of new user equipment (UE) on the network based on network data analytics in accordance with embodiments described herein;

10           Figure 6 illustrates a logical flow diagram showing one embodiment of a process for selecting the UPF, which is useful in the process of Figure 5 in accordance with embodiments described herein;

15           Figure 7 illustrates a logical flow diagram showing one embodiment of a process for selecting the UPF using artificial intelligence (AI) or machine learning (ML) algorithms to perform predictive analysis of throughput, which is useful in the process of Figure 6 in accordance with embodiments described herein;

            Figure 8 illustrates a logical flow diagram showing one embodiment of a process for UPF load balancing based on special considerations for low latency traffic in accordance with embodiments described herein;

20           Figure 9 illustrates a logical flow diagram showing one embodiment of a process for selecting the UPF based on the location of the new UE and load-regions for each UPF defined by load thresholds for non-low latency traffic, which is useful in the process of Figure 8 in accordance with embodiments described herein;

25           Figure 10 illustrates a logical flow diagram showing one embodiment of a process for selecting the UPF based on whether the network traffic is identified as low latency, which is useful in the process of Figure 9 in accordance with embodiments described herein;

30           Figure 11 illustrates a logical flow diagram showing one embodiment of a process for UPF load balancing supporting multiple slices, maintaining several load-thresholds for each UPF and each slice depending on the UPF and network slice capacity in accordance with embodiments described herein;

            Figure 12 illustrates a logical flow diagram showing one embodiment of a process for selecting the network slice based on generated weights, which is useful in the process of Figure 11 in accordance with embodiments described herein;

Figure 13 illustrates a logical flow diagram showing one embodiment of a process for selecting the network slice based on determined load-regions for each slice and weights generated based on the determined load regions, which is useful in the process of Figure 12 in accordance with embodiments described herein;

5           Figure 14 illustrates a logical flow diagram showing one embodiment of a process for UPF load balancing using predicted central processing unit (CPU) utilization and/or predicted memory utilization of new user equipment (UE) on the network based on network data analytics in accordance with embodiments described herein;

10           Figure 15 illustrates a logical flow diagram showing one embodiment of a process for selecting the UPF, which is useful in the process of Figure 14 in accordance with embodiments described herein;

            Figure 16 illustrates a logical flow diagram showing one embodiment of a process for selecting the UPF using AI or machine learning (ML) algorithms to  
15           perform predictive analysis of CPU utilization and/or predicted memory utilization, which is useful in the process of Figure 15 in accordance with embodiments described herein;

            Figure 17 illustrates a chart showing an example of possible load-regions that a current load of a UPF may be determined to fall within that may be used in the  
20           processes of Figures 2-7 and 15-17 in accordance with embodiments described herein.

            Figure 18 illustrates a timeline showing an example of possible UPF load balancing that may occur according to the processes of Figures 2-4 with two UPFs determined to fall within particular different load-regions of those shown in Figure 17  
25           at particular times in accordance with embodiments described herein;

            Figure 19 illustrates a timeline showing an example of possible UPF load balancing that may occur according to the processes of Figures 5-7 with two UPFs determined to fall within particular different load-regions of those shown in Figure 17 at particular times in accordance with embodiments described herein.

30           Figure 20 illustrates a chart showing an example of possible load-regions for non-low latency traffic that a current load of a UPF may be determined to fall within that may be used in the processes of Figures 8-10 in accordance with embodiments described herein.

Figure 21 shows a system diagram that describe various implementations of computing systems for implementing embodiments described herein.

#### DETAILED DESCRIPTION

The following description, along with the accompanying drawings, sets forth  
5 certain specific details in order to provide a thorough understanding of various  
disclosed embodiments. However, one skilled in the relevant art will recognize that  
the disclosed embodiments may be practiced in various combinations, without one or  
more of these specific details, or with other methods, components, devices,  
materials, etc. In other instances, well-known structures or components that are  
10 associated with the environment of the present disclosure, including but not limited to  
the communication systems and networks, have not been shown or described in  
order to avoid unnecessarily obscuring descriptions of the embodiments.  
Additionally, the various embodiments may be methods, systems, media, or devices.  
Accordingly, the various embodiments may be entirely hardware embodiments,  
15 entirely software embodiments, or embodiments combining software and hardware  
aspects.

Throughout the specification, claims, and drawings, the following terms take  
the meaning explicitly associated herein, unless the context clearly dictates  
otherwise. The term “herein” refers to the specification, claims, and drawings  
20 associated with the current application. The phrases “in one embodiment,” “in  
another embodiment,” “in various embodiments,” “in some embodiments,” “in other  
embodiments,” and other variations thereof refer to one or more features, structures,  
functions, limitations, or characteristics of the present disclosure, and are not limited  
to the same or different embodiments unless the context clearly dictates otherwise.  
25 As used herein, the term “or” is an inclusive “or” operator, and is equivalent to the  
phrases “A or B, or both” or “A or B or C, or any combination thereof,” and lists with  
additional elements are similarly treated. The term “based on” is not exclusive and  
allows for being based on additional features, functions, aspects, or limitations not  
described, unless the context clearly dictates otherwise. In addition, throughout the  
30 specification, the meaning of “a,” “an,” and “the” include singular and plural  
references.

Figure 1 illustrates a context diagram of an environment 100 in which UPF load balancing may be implemented in accordance with embodiments described herein.

5 UEs 110, such as cellular telephones or other Internet-of-Things (IoT) devices use 5G wireless cellular telecommunication technology defined by standards set by 3GPP and International Telecommunications Union (ITU) to get data connectivity between applications on the UE and Data Networks (DNs) such as the Internet or private corporate networks. Almost all applications running on the UE, including voice, require such data connectivity. A Protocol Data Unit (PDU) session provides  
10 connectivity between applications on a UE and a DN. The UE receives services through a PDU session, which is a logical connection between the UE and DN. A DN is identified by a Data Network Name (DNN). PDU sessions can provide different types of transport services corresponding to the nature of the PDU(s) carried over the PDU session. In various embodiments, a PDU session may be  
15 associated with a single DNN and with a single slice identified by Single-Network Slice Selection Assistance Information (S-NSSAI).

The UPF is one of the network functions (NFs) of the 5GC. The UPF, comprising UPF1 104 and UPF2 106 in the present example, is responsible for packet routing and forwarding, packet inspection, quality of service (QoS) handling,  
20 and interconnecting external PDU sessions with the DN. Although two UPFs (UPF1 104 and UPF2 106) are shown in the present example, additional UPFs may be utilized in various other embodiments. Each UPF (e.g., UPF1 104 and UPF2 106) is a virtual network function responsible for PDU sessions between the UEs 110 and the DN by anchoring the PDU sessions of various UEs 110 on the individual UPF.  
25 The SMF 102 is also one of the NFs of the 5GC and is primarily responsible for interacting with the decoupled data plane, creating updating and removing PDU sessions, selecting particular UPFs on which to anchor PDU sessions when new UEs appear on the network and managing session context with the UPF. Many of such functions are described in the 3GPP TS 23.501 specification.

30 A network function, such as the SMF 102 and the UPF, such as UPF1 104 and UPF2 106, can be implemented either as a network elements on dedicated hardware, as a software instance running on dedicated hardware, or as a virtualized function instantiated on an appropriate platform, e.g., a cloud infrastructure. In the present example, UPF1 104 is implemented at data center 1 and UPF2 106 is

implemented at data center 2, which is geographically separated from data center 1. The SMF 102 sends messages to the UPF (comprising UPF 1 104 and UPF 2 in the present example) over the N4 reference interface using the Packet Forwarding Control Protocol (PFCP). The PFCP may employ UDP port (8805) and is defined to support Control and User Plane Separation (CUPS). Decoupling other control plane functions from the user plane, together with the 5G Core Access and Mobility Management Function (AMF) (not shown), the SMF 102 performs the role of Dynamic Host Control Protocol (DHCP) server and Internet Protocol (IP) Address Management (IPAM) system. Together with the UPF, the SMF 102 maintains a record of PDU session state by means of a 24bit PDU Session ID. The SMF 102 sets configuration parameters in the UPF that define traffic steering parameters and ensure the appropriate routing of packets while guaranteeing the delivery of incoming packets, though a Downlink (DL) data notification.

In the present example embodiment, each UPF1 104 and UPF2 106 may have the ability to establish network connectivity and anchor PDU sessions of any UE on the network via various cellular telecommunication base stations and associated antennas 108. To maximize network performance, PDU sessions are by default anchored on the UPF at the data center that is closest geographically to the UE, as illustrated by most of the dashed lines in Figure 1 for UEs 110 (and an operator defines a service area for each UPF). However, each UPF (e.g., UPF1 104 and UPF2 106) has a maximum network capacity to handle PDU sessions anchored thereon and the associated network traffic. Thus, PDU sessions anchored on a particular UPF (e.g., UPF1 104) and their associated network traffic may cause the UPF to near its maximum capacity or become overloaded. UPF load balancing may then cause the PDU session of the next new UE appearing on the network (e.g., UE 112) to be anchored on a UPF at a data center (e.g., UPF2 106) that is further away than the data center that is closest geographically to the UE. In the present example, UPF1 104 is at or near its maximum capacity with the PDU sessions of all the other UEs currently anchored on it, so UE 112 has a PDU session anchored on UPF2 106 (as shown by dashed line 114) instead of UPF1 104, even though data center 2 of UPF2 106 is further away from the UE 112 than data center 1 of UPF1 104. In various embodiments described herein, there are different particular scenarios and rules in which UPF load balancing may cause the PDU session of the next new UE appearing on the network to be anchored on a UPF at a data center



that is further away than the data center that is closest geographically to the UE, which improves overall UPF load balancing and network performance.

Figure 2 illustrates a logical flow diagram showing one embodiment of a process 200 for load balancing based on current UPF load and thresholds that  
5 depend on UPF capacity in accordance with embodiments described herein.

At 202, the SMF 102 maintains load thresholds for each user plane function (UPF) of a plurality of UPFs in a cellular telecommunication network. The plurality of UPFs serve as anchor points between UE in the cellular telecommunication network and a DN. Each UPF of the plurality of UPFs is a virtual network function  
10 responsible for interconnecting PDU sessions between the UE and the DN by anchoring the PDU sessions on individual UPFs. The load thresholds for each UPF depend on a respective capacity of each UPF to have PDU sessions anchored thereon. In the present example embodiment, an amount of load put on a UPF by a UE appearing in the cellular telecommunication network is assumed to be identical  
15 for all UEs appearing in the cellular telecommunication network.

At 204, the SMF 102 receives a request to anchor on a UPF a PDU session of a new UE newly appearing on the cellular telecommunication network.

At 206, the SMF 102 selects a UPF of the plurality of UPFs on which to anchor the PDU session based on a location of the new UE and determined load-  
20 regions for each UPF of the plurality of UPFs defined by the load thresholds.

At 208, the SMF 102 anchors the PDU session of the new UE to the selected UPF.

Figure 3 illustrates a logical flow diagram showing one embodiment of a process 300 for selecting the UPF based on generated weights, which is useful in  
25 the process 200 of Figure 2 in accordance with embodiments described herein.

At 302, the SMF 102 generates weights for selecting the UPF based on the determined load-regions.

At 304, the SMF 102 selects the UPF based on the generated weights.

Figure 4 illustrates a logical flow diagram showing one embodiment of a  
30 process 400 for selecting the UPF based on the determined load-regions for a plurality of UPFs and the weights generated based on the determined load regions, which is useful in the process 300 of Figure 3 in accordance with embodiments described herein.

At 402, the SMF 102 generates multiple load-regions. Each load-region corresponds to a different range of current load of a UPF defined by one or more of lower and upper threshold percentages of UPF load capacity.

At 404, the SMF 102 receives the request to anchor the PDU session.

5 At 406, the SMF 102 determines a load region from the multiple load-regions that a current load of the UPF falls within.

At 408, the SMF 102 determines whether there are additional UPFs in the plurality of UPFs on which the PDU session may be anchored. If it is determined there are additional UPFs on which the PDU session may be anchored, then the  
10 process 400 proceeds back to 406 to determine a load region from the multiple load-regions that a current load of the additional UPF falls within. If it is determined there are not additional UPFs on which the PDU session may be anchored, then the process 400 proceeds to 410.

At 410, the SMF 102 selects a UPF of the plurality of UPFs based on the  
15 determined load-regions for the plurality of UPFs and the weights generated based on the determined load regions.

In an example embodiment, the SMF 102 generates a lowest load-region indicating a current UPF load less than a first threshold percentage of UPF capacity; generates one or more intermediate non-overlapping load-regions each defined by  
20 respective lower and upper threshold percentages of UPF capacity and indicating a current load greater than the lowest load-region; and generates a highest load-region indicating a current UPF load greater than a second threshold percentage of UPF capacity and greater than the intermediate non-overlapping load-regions.

In the present example embodiment, each UPF is associated with a different  
25 respective geographic UPF service area. Selecting the UPF based on the generated weights and the determined load-regions for the UPFs may include determining a particular UPF has (i.e., is at a data center in) a respective geographic area within which the location of the new UE falls (i.e., is at a data center that is closest geographically to the UE compared to data centers of other UPFs). The particular  
30 UPF may then be selected in response to the determined load-region of the particular UPF being a load-region indicating its current load is below a threshold capacity.

In some embodiments, the SMF 102 determines a particular UPF has a respective geographic area within which the location of the new UE falls. The SMF

102 determines whether the particular UPF has a determined load region indicating its current load is in a different load region indicating a higher current load of the particular UPF than a current load of another UPF. In response to this determination, the SMF 102 weights the selection of a UPF. In particular, the UPF selection by the SMF for load-balancing is based on weighted scheduling of load (UEs) on the UPFs. This weighted scheduling may be credit/token-based (e.g., weighted round robin) or probability based (e.g., using statistical based scheduling algorithms using probability). For example, the SMF 102 may weight the selection of a UPF such that a probability that the particular UPF is selected is lower than a probability of selection of the other UPF. In some embodiments, the selection of the UPF is weighted by using credit/token-based weighted scheduling or probability-based weighted scheduling such that the frequency of selection of the particular UPF decreases as a difference between a higher current load of the particular UPF and a lower current load of at least one UPF of the plurality of UPFs increases, as indicated by the load regions determined for each UPF of the plurality of UPFs..

Figure 5 illustrates a logical flow diagram showing one embodiment of a process 500 for UPF load balancing using predicted throughput of a new UE on the network based on network data analytics in accordance with embodiments described herein.

At 502, the SMF 102 maintains load thresholds for each user plane function (UPF) of a plurality of UPFs in a cellular telecommunication network. The plurality of UPFs serve as anchor points between UE in the cellular telecommunication network and a DN. Each UPF of the plurality of UPFs is a virtual network function responsible for interconnecting PDU sessions between the UE and the DN by anchoring the PDU sessions on individual UPFs. The load thresholds for each UPF depend on a respective capacity of each UPF to have PDU sessions anchored thereon. However, in the present example embodiment, an amount of load put on a UPF by a UE appearing in the cellular telecommunication network is not assumed to be identical for all UEs appearing in the cellular telecommunication network.

At 504, the SMF 102 receives a request to anchor on a UPF a PDU session of a new UE newly appearing on the cellular telecommunication network.

At 506, the SMF 102 selects a UPF of the plurality of UPFs on which to anchor the PDU session based on a location of the new UE, determined load-regions for each UPF of the plurality of UPFs defined by the load thresholds and

predicted throughput of the new UE based on network data analytics. In an example embodiment, the network data analytics is provided via a network data analytics function (NWDAF) of a 5G mobile network of which the cellular telecommunication network is comprised.

5           At 508, the SMF 102 anchors the PDU session of the new UE to the selected UPF.

Figure 6 illustrates a logical flow diagram showing one embodiment of a process 600 for selecting the UPF, which is useful in the process 500 of Figure 5 in accordance with embodiments described herein.

10           At 602, in selecting the UPF, the SMF 102 uses the network data analytics to predict throughput of the UE and load on a UPF of the new UE appearing on the cellular telecommunication network based on the predicted throughput.

At 604, the SMF 102 selects a UPF of the plurality of UPFs on which to anchor the PDU session based on a location of the new UE, load-regions for each UPF of the plurality of UPFs defined by the load thresholds and the predicted load of  
15           the new UE on a UPF.

Figure 7 illustrates a logical flow diagram showing one embodiment of a process 700 for selecting the UPF using artificial intelligence (AI) or machine learning (ML) algorithms to perform predictive analysis of throughput, which is useful  
20           in the process 600 of Figure 6 in accordance with embodiments described herein.

At 702, in using the network data analytics to predict throughput of the new UE and load on a UPF, the SMF 102 uses artificial intelligence (AI) or machine learning (ML) algorithms to perform predictive analysis of throughput of the new UE and resulting load on a UPF of the new UE appearing on the cellular  
25           telecommunication network based on historical activity of the new UE appearing on the cellular telecommunication network.

At 704, the SMF 102 implements a weighted scheduling of load on UPFs to achieve UPF load-balancing based on the predicted throughput of the new UE and resulting predicted load on a UPF of the new UE. This weighted scheduling can be  
30           implemented using credit/token based scheduling algorithms or statistical based scheduling algorithms (using probability). For example, in one embodiment, the SMF 102 changes a probability of whether a particular UPF of the plurality of UPFs will be selected based on the predicted throughput of the new UE and resulting predicted load on a UPF of the new UE. In an example embodiment, the SMF 102

changes a probability of whether a particular UPF will be selected based on the predicted throughput of the new UE and resulting predicted load on a UPF of the new UE. For example, the SMF 102 may increase a probability that a particular UPF will be selected in response to a current load of the particular UPF being currently in a particular load-region as compared to other UPFs and the predicted load being at a particular level. The SMF 102 may decrease a probability that the particular UPF will be overloaded beyond a threshold amount compared to other UPFs based on the predicted load by changing the probability of whether the particular UPF will be selected based on the predicted load. Such load balancing may instead be achieved using credit/token based scheduling (e.g., weighted round robin).

Figure 8 illustrates a logical flow diagram showing one embodiment of a process 800 for UPF load balancing based on special considerations for low latency traffic in accordance with embodiments described herein.

At 802, the SMF 102 maintains load thresholds for each user plane function (UPF) of a plurality of UPFs in a cellular telecommunication network. The plurality of UPFs serve as anchor points between UE in the cellular telecommunication network and a DN. Each UPF of the plurality of UPFs is a virtual network function responsible for interconnecting PDU sessions between the UE and the DN by anchoring the PDU sessions on individual UPFs. The load thresholds for each UPF depend on a respective capacity of each UPF to have PDU sessions anchored thereon. In the present example embodiment, an amount of load put on a UPF by a UE appearing in the cellular telecommunication network is assumed to be identical for all UEs appearing in the cellular telecommunication network.

In some embodiments, the load thresholds may be reduced by a percentage amount of capacity dedicated for low-latency network traffic. For example, a percentage amount of capacity dedicated for low-latency network traffic may be 10% and thus the load thresholds for non-low latency traffic (such as the thresholds maintained in the process 200 of Figure 2) may be reduced by 10%.

At 804, the SMF 102 receives a request to anchor on a UPF a PDU session of a new UE newly appearing on the cellular telecommunication network.

At 806, the SMF 102 selects a UPF of the plurality of UPFs on which to anchor the PDU session based on whether traffic of the PDU session is identified as low latency and a location of the new UE.

At 808, the SMF 102 anchors the PDU session of the new UE to the selected UPF.

Figure 9 illustrates a logical flow diagram showing one embodiment of a process 900 for selecting the UPF based on the location of the new UE and load-  
5 regions for each UPF defined by load thresholds for non-low latency traffic, which is useful in the process 800 of Figure 8 in accordance with embodiments described herein.

At 902, the SMF 102 receives a request to anchor on a UPF a PDU session of a new UE newly appearing on the cellular telecommunication network.

10 At 904, the SMF 102 determines whether the traffic of the PDU session is identified as low latency. In the present example embodiment, the selection of the UPF is based on dedicating a percentage of capacity of each UPF of the plurality of UPFs to low-latency traffic of PDU sessions. Latency may be measured in the time elapsed from when the client sends the first byte of a request to the moment the  
15 server receives it, or it may be measured by the total journey time for a packet to travel to the server and then back to the client. In the present example, on the downlink, the latency is measured from the time that the UPF receives the packet until the time that the packet is delivered to the UE. On the uplink, the latency is measured from the time that the UE sends the packet until the time that the packet is  
20 received by the UPF. For example, low latency network traffic may support operations that require near real-time access to rapidly changing data. Low latency is desirable in a wide range of use cases. In a general sense, lower latency is nearly always an improvement over slower packet transport. Low latency is desirable in online gaming as it contributes to a more realistic gaming environment. The term  
25 low latency is often used to describe specific business use cases, in particular high-frequency trading in capital markets. If traffic of the PDU session is identified as low latency, then the process 900 proceeds to 906. If traffic of the PDU session is not identified as low latency, then the process 900 proceeds to 908.

At 906, the SMF 102 selects a UPF having a closest associated location to a  
30 current location of the new UE.

At 908, the SMF selects a UPF based on the location of the new UE and load- regions for each UPF of the plurality of UPFs defined by the load thresholds for non- low latency traffic.

Figure 10 illustrates a logical flow diagram showing one embodiment of a process 1000 for selecting the UPF based on whether the network traffic is identified as low latency, which is useful in the process 900 of Figure 9 in accordance with embodiments described herein.

5           At 1002, the SMF 102 receives a request to anchor on a UPF a PDU session of a new UE newly appearing on the cellular telecommunication network.

          At 1004, the SMF 102 determines whether the traffic of the PDU session is identified as low latency. If traffic of the PDU session is identified as low latency, then the process 1000 proceeds to 1006. If traffic of the PDU session is not  
10 identified as low latency, then the process 1000 proceeds to 1008.

          At 1006, the SMF 102 selects a UPF having a closest associated location to a current location of the new UE.

          At 1008, the SMF 102 selects a UPF based on weights for selecting the UPF generated based on the load-regions, wherein each load-region corresponds to a  
15 different range of current load of the UPF defined by one or more of lower and upper threshold percentages of load capacity of the UPF.

          In some embodiments, if each UPF of the plurality of UPFs is identified as currently having a current load falling within a low load-region defined by a current load below a particular threshold, then the SMF 102 selects a UPF having a closest  
20 associated location to a current location of the new UE.

Figure 11 illustrates a logical flow diagram showing one embodiment of a process 1100 for UPF load balancing supporting multiple slices, maintaining several load-thresholds for each UPF and each slice depending on the UPF and network slice capacity in accordance with embodiments described herein.

25           At 1102, the SMF 102 maintains load thresholds for each network slice of a plurality of network slices. In the present example embodiment, each network slice of each respective set of network slices comprises of a set of virtual network resources and network traffic flows associated with the network slice and represents an independent virtualized instance of a network defined by allocation of a subset of  
30 available network resources in the cellular telecommunication network. Each network slice of the plurality of network slices is supported by a respective user plane function (UPF) of a plurality of UPFs in a cellular telecommunication network. The plurality of UPFs serve as anchor points between user equipment (UE) in the cellular telecommunication network and a data network (DN). Each UPF of the

plurality of UPFs is a virtual network function responsible for interconnecting packet data unit (PDU) sessions between the user equipment (UE) and the DN by anchoring the PDU sessions on individual UPFs. The load thresholds for each network slice depend on a respective capacity of each network slice and total  
5 capacity of each UPF supporting each network slice to have PDU sessions anchored thereon. An amount of load put on a network slice by a UE appearing in the cellular telecommunication network is assumed to be identical for all UEs appearing in the cellular telecommunication network.

At 1104, the SMF 102 receives a request to anchor on a UPF a PDU session  
10 of a new UE newly appearing on the cellular telecommunication network.

At 1106, the SMF 102 selects a network slice of the plurality of network slices on which to anchor the PDU session based on a location of the new UE and determined load-regions for each network slice of the plurality of network slices defined by the load thresholds.

At 1108, the SMF 102 anchors the PDU session of the new UE to the  
15 selected network slice and the respective UPF supporting the selected network slice.

Figure 12 illustrates a logical flow diagram showing one embodiment of a process 1200 for selecting the network slice based on generated weights, which is useful in the process 1100 of Figure 11 in accordance with embodiments described  
20 herein.

At 1202, the SMF 102 generates weights for selecting the network slice based on the determined load-regions.

At 1204, the SMF 102 selects the network slice based on the generated weights.

Figure 13 illustrates a logical flow diagram showing one embodiment of a process 1300 for selecting the network slice based on determined load-regions for each slice and weights generated based on the determined load regions, which is useful in the process 1200 of Figure 12 in accordance with embodiments described  
25 herein.

At 1302, the SMF 102 generates multiple load-regions. Each load-region corresponds to a different range of current load of a network slice defined by one or more of lower and upper threshold percentages of network slice load capacity. For example, in one embodiment, the SMF 102 may generate a lowest load-region indicating a current network slice load less than a first threshold percentage of  
30



network slice capacity; generate one or more intermediate non-overlapping load-regions each defined by respective lower and upper threshold percentages of network slice capacity and indicating a current load greater than the lowest load-region; and generate a highest load-region indicating a current network slice load  
5 greater than a second threshold percentage of network slice capacity and greater than the intermediate load-region(s).

At 1304, the SMF 102 receives the request to anchor the PDU session.

At 1306, the SMF 102, in response to receiving the request to anchor the PDU session, determines a load region from the multiple load-regions that a current  
10 load of a network slice falls within.

At 1306, the SMF 102 determines whether there are any additional network slides on which the PDU session may be anchored. If it is determined there are additional network slides on which the PDU session may be anchored, then process 1300 proceeds back to 1306 to determine a load region from the multiple load-  
15 regions that a current load of the additional network slice falls within. If it is determined there are not additional network slides on which the PDU session may be anchored, then the process 1300 proceeds to 1310.

At 1310, the SMF 102 selects a network slice of the plurality of network slices based on the determined load-regions for the plurality of network slices and the  
20 weights generated based on the determined load regions. In some embodiments, each network slice of the plurality of network slices is associated with a respective geographic area of the respective UPF supporting the network slice (i.e., the geographic area of the data center of the UPF). Selecting a network slice based on the generated weights and the determined load-regions for the plurality of network  
25 slices may include determining a particular network slice of the plurality of network slices is associated with a respective geographic area within which the location of the new UE falls. The SMF 102 may then select the particular network slice in response the determined load-region of the particular network slice being a load-region indicating a current load of the particular network slice is below a threshold  
30 capacity.

In some embodiments, selecting a network slice based on the generated weights and the determined load-regions for the plurality of network slices may include determining whether the particular network slice has a determined load region indicating a current load of the particular network slice is in a different load

region indicating a higher current load of the particular network slice than a current load of another network slice. In response to this, the SMF 102 may weight the selection of a network slice of the plurality of network slices such that the particular network slice is not overloaded compared to the other network slice by using  
5 credit/token-based weighted scheduling or probability-based weighted scheduling.

In some embodiments, the selection of a network slice includes determining a particular network slice of the plurality of network slices is associated with a respective geographic area within which the location of the new UE falls. The selection of the network slice is then weighted by using credit/token-based weighted  
10 scheduling or probability-based weighted scheduling such that the frequency of selection of the particular network slice decreases as a difference between a higher current load of the particular network slice and a lower current load of at least one network slice of the plurality of network slices increases, as indicated by the load regions determined for each network slice of the plurality of network slices.

15 Figure 14 illustrates a logical flow diagram showing one embodiment of a process 1400 for UPF load balancing using predicted CPU utilization and/or predicted memory utilization of new UE on the network based on network data analytics in accordance with embodiments described herein.

At 1402, the SMF 102 maintains load thresholds for each UPF of a plurality of  
20 UPFs in a cellular telecommunication network. The plurality of UPFs serve as anchor points between UE in the cellular telecommunication network and a DN. Each UPF of the plurality of UPFs is a virtual network function responsible for interconnecting PDU sessions between the UE and the DN by anchoring the PDU sessions on individual UPFs. The load thresholds for each UPF depend on a  
25 respective capacity of each UPF to have PDU sessions anchored thereon.

At 1404, the SMF 102 receives a request to anchor on a UPF a PDU session of a new UE newly appearing on the cellular telecommunication network.

At 1406, the SMF 102 selects a UPF of the plurality of UPFs on which to anchor the PDU session based on a location of the new UE, load-regions for each  
30 UPF of the plurality of UPFs defined by the load thresholds and one or more of predicted CPU utilization and predicted memory utilization of the new UE based on network data analytics. In some embodiments, the network data analytics may be provided via an NWDAF of a 5G mobile network of which the cellular telecommunication network is comprised.

Figure 15 illustrates a logical flow diagram showing one embodiment of a process 1500 for selecting the UPF, which is useful in the process 1400 of Figure 14 in accordance with embodiments described herein.

At 1502, the SMF 102 uses the network data analytics to predict one or more  
5 of CPU utilization and memory utilization of the UE and load on a UPF of the new UE appearing on the cellular telecommunication network based on one or more of the predicted CPU utilization and predicted memory utilization.

At 1504, the SMF 102 selects a UPF of the plurality of UPFs on which to  
10 anchor the PDU session based on a location of the new UE, load-regions for each UPF of the plurality of UPFs defined by the load thresholds and the predicted load of the new UE on a UPF.

Figure 16 illustrates a logical flow diagram showing one embodiment of a process for 1600 selecting the UPF using AI or machine learning ML algorithms to perform predictive analysis of CPU utilization and/or predicted memory utilization,  
15 which is useful in the process 1500 of Figure 15 in accordance with embodiments described herein.

At 1602, the SMF 102 uses artificial intelligence (AI) or machine learning (ML)  
algorithms to perform predictive analysis of one or more of CPU utilization and  
memory utilization of the new UE and resulting load on a UPF of the new UE  
20 appearing on the cellular telecommunication network based on historical activity of the new UE appearing on the cellular telecommunication network.

At 1604, the SMF 102 implements a weighted scheduling of load on UPFs to  
achieve UPF load-balancing based on one or more of the predicted CPU utilization  
and predicted memory utilization of the new UE and resulting predicted load on a  
25 UPF of the new UE. This weighted scheduling can be implemented using credit/token based scheduling algorithms or statistical based scheduling algorithms (using probability). For example, in one embodiment, the SMF 102 may increase a probability that a particular UPF will be selected in response to a current load of the particular UPF being currently in a particular load-region as compared to other UPFs  
30 of the plurality of UPFs and the resulting predicted load resulting from one or more of the predicted CPU utilization and predicted memory utilization being at a particular level. The SMF 102 may decrease a probability that the particular UPF in the plurality of UPFs will be overloaded beyond a threshold amount compared to other UPFs in the plurality of UPFs based on the resulting predicted load by changing the

probability of whether the particular UPF will be selected based on the resulting predicted load resulting from one or more of the predicted CPU utilization and predicted memory utilization. Such load balancing may instead be achieved using credit/token based scheduling (e.g., weighted round robin).

5           Figure 17 illustrates a chart 1700 showing an example of possible load-regions that a current load of a UPF may be determined to fall within that may be used in the processes of Figures 2-7 and 15-17 in accordance with embodiments described herein.

10           The chart 1700 indicates the load region 1702 and the range of percentage of total load capacity 1704 for each load-region 1702 defined by respective load thresholds indicated in the chart 1700. In the present example, the “Low” load region indicates a current UPF load of less than 30% of the total load capacity. The “Medium” load region indicates a current UPF load of greater than or equal to 30% and less than 55% of the total load capacity. The “High” load region indicates a  
15           current UPF load of greater than or equal to 55% and less than 75% of the total load capacity. The “Very High” load region indicates a current UPF load of greater than or equal to 75% of the total load capacity. There may be additional or different load regions in various other embodiments. The chart 1700 may be stored or represented as a data structure in computer memory and may be maintained and/or  
20           accessible by the SMF 102.

          Figure 18 illustrates a timeline 1800 showing an example of possible UPF load balancing that may occur according to the processes of Figures 2-4 with two UPFs determined to fall within particular different load-regions of those shown in the chart 1700 Figure 17 at particular times in accordance with embodiments described  
25           herein.

          For each example point in time  $t=1$ ,  $t=2$ ,  $t=3$  and  $t=4$ , shown in the timeline 1800 is the determined UPF1 load 1802, the determined UPF2 load 1806 and the corresponding UPF1 load balancing action 1804 and UPF2 load balancing action 1808.

30           Referring now also to Figure 1, in the present example, at time  $t=1$ , the SMF 102 determines that the current load of UPF1 104 is 20% of the total load capacity for UPF1 104 and thus determines the load region or UPF1 104 is “low” based on the load thresholds in the chart 1700 of Figure 17. At time  $t=1$ , the SMF 102 also determines that the current load of UPF2 106 is 2% of the total load capacity for

UPF2 106 and thus determines the load region or UPF2 106 is also “low” based on the load thresholds in the chart 1700 of Figure 17. Based on the determined load regions for UPF1 104 and UPF2 106, any new UE in the region of UPF1 104 (e.g., the region of UPF1 104 being a respective geographic area associated with data center 1) are anchored on UPF1 104 and any new UE in the region of UPF2 106 (e.g., the region of UPF2 106 being a respective geographic area associated with data center 2) are anchored on UPF2 106.

At time  $t=2$ , the SMF 102 determines that the current load of UPF1 104 is 45% of the total load capacity for UPF1 104 and thus determines the load region or UPF1 104 is “medium” based on the load thresholds in the chart 1700 of Figure 17. At time  $t=2$ , the SMF 102 also determines that the current load of UPF2 106 is 10% of the total load capacity for UPF2 106 and thus determines the load region or UPF2 106 is “low” based on the load thresholds in the chart 1700 of Figure 17. Based on the determined load regions for UPF1 104 and UPF2 106, the weights determined by the SMF 102 for UPF1 104 load balancing are 1 & 2. In particular, for every three new UEs in the region of UPF1 104 (data center 1), one is anchored on UPF1 104 and the other two are anchored on UPF2 106. In various embodiments, the load balancing can be done using credit/token based scheduling algorithms (using weighted round robin) or done randomly (using statistical based scheduling algorithms using probability): with probability  $1/3$  a new UE is anchored on UPF1 104 and with probability  $2/3$  the new UE is anchored on UPF2 106. Any new UE in the region of UPF2 106 are anchored on UPF2 106 (data center 2).

At time  $t=3$ , the SMF 102 determines that the current load of UPF1 104 is 65% of the total load capacity for UPF1 104 and thus determines the load region or UPF1 104 is “high” based on the load thresholds in the chart 1700 of Figure 17. At time  $t=3$ , the SMF 102 also determines that the current load of UPF2 106 is 19% of the total load capacity for UPF2 106 and thus determines the load region or UPF2 106 is “low” based on the load thresholds in the chart 1700 of Figure 17. Based on the determined load regions for UPF1 104 and UPF2 106, the weights determined by the SMF 102 for UPF1 104 load balancing are 1 & 4. In particular for every five new UEs in the region of UPF1 (data center1), one is anchored on UPF1 and the other four are anchored on UPF2. In various embodiments, the load balancing can be done by the SMF 102 using weighted round robin or done randomly, with probability  $1/5$  a new UE is anchored on UPF1 and with probability  $4/5$  the new UE is anchored

on UPF2). Any new UE in the region of UPF2 106 (e.g., the region of UPF2 106 being a respective geographic area associated with data center 2) are anchored on UPF2 106 (data center 2).

At time  $t=4$ , the SMF 102 determines that the current load of UPF1 104 is  
5 75% of the total load capacity for UPF1 104 and thus determines the load region or UPF1 104 is “very high” based on the load thresholds in the chart 1700 of Figure 17. At time  $t=4$ , the SMF 102 also determines that the current load of UPF2 106 is 25% of the total load capacity for UPF2 106 and thus determines the load region or UPF2 106 is “low” based on the load thresholds in the chart 1700 of Figure 17. Based on  
10 the determined load regions for UPF1 104 and UPF2 106, any new UE in the region of UPF1 104 (data center 1) are anchored on UPF2 106 and any new UE in the region of UPF2 106 (data center 2) are also anchored on UPF2 106. In some embodiments, the credit based scheduling method, such as the weighted round robin scheduling, disclosed herein is based on a probabilistic model. Alternatively, a  
15 credit/token based weighted round robin scheduling can be used in various other embodiments.

Figure 19 illustrates a timeline showing an example of possible UPF load balancing that may occur according to the processes of Figures 5-7 with two UPFs determined to fall within particular different load-regions of those shown in Figure 17  
20 at particular times in accordance with embodiments described herein.

For each example point in time  $t=1$ ,  $t=2$ ,  $t=3$  and  $t=4$ , shown in the timeline 1900 is the determined UPF1 load 1902, the determined UPF2 load 1906, the corresponding UPF1 load balancing action 1904 and UPF2 load balancing action 1908, as well as example new UE predicted loads 1910. For example, the predicted  
25 load may be based on data from the network data analytics function (NWDAF). The predicted load may be measured in units based on throughput (e.g., packets per second, bytes per second, and/or bits per second), CPU utilization (e.g., CPU clock cycles, clock ticks, CPU time, CPU time per second, process time, percentage of CPU capacity utilization) and/or memory utilization, (megabytes of memory, and/or  
30 percentage of memory capacity utilization) or any combination thereof.

Referring now also to Figure 1, in the present example, at time  $t=1$ , the SMF 102 determines that the current load of UPF1 104 is 20% of the total load capacity for UPF1 104 and thus determines the load region or UPF1 104 is “low” based on the load thresholds in the chart 1700 of Figure 17. At time  $t=1$ , the SMF 102 also

determines that the current load of UPF2 106 is 2% of the total load capacity for UPF2 106 and thus determines the load region or UPF2 106 is also “low” based on the load thresholds in the chart 1700 of Figure 17. In this case, based on the determined load regions for UPF1 104 and UPF2 106, regardless of any predicted  
5 load of the new UE, any new UE in the region of UPF1 104 (data center 1) are anchored on UPF1 104 and any new UE in the region of UPF2 106 (data center 2) are anchored on UPF2 106.

At time  $t=2$ , the SMF 102 determines that the current load of UPF1 104 is 45% of the total load capacity for UPF1 104 and thus determines the load region or  
10 UPF1 104 is “medium” based on the load thresholds in the chart 1700 of Figure 17. At time  $t=2$ , the SMF 102 also determines that the current load of UPF2 106 is 10% of the total load capacity for UPF2 106, and thus determines the load region or UPF2 106 is “low” based on the load thresholds in the chart 1700 of Figure 17. Based on the determined load regions for UPF1 104 and UPF2 106, the weights determined by  
15 the SMF 102 for UPF1 104 load balancing are 1 & 2. For example, based on the detected predicted load of the new UE appearing on the network, if the SMF 102 determines that the predicted load of the new UE is 2 units of load, the SMF 102 will perform load balancing such that there is a 43% probability the SMF 102 anchors the UE to UPF1 104 and a 57% probability the SMF 102 anchors the UE to UPF2 106.  
20 If the SMF 102 determines that the predicted load of the new UE is 1 unit of load, the SMF 102 will perform load balancing such that there is a 25% probability the SMF 102 anchors the UE to UPF1 104 and a 75% probability the SMF 102 anchors the UE to UPF2 106. Any new UE in the region of UPF2 (data center 2) 106 are anchored on UPF2 106.

25 At time  $t=3$ , the SMF 102 determines that the current load of UPF1 104 is 65% of the total load capacity for UPF1 104 and thus determines the load region or UPF1 104 is “high” based on the load thresholds in the chart 1700 of Figure 17. At time  $t=3$ , the SMF 102 also determines that the current load of UPF2 106 is 19% of the total load capacity for UPF2 106 and thus determines the load region or UPF2  
30 106 is “low” based on the load thresholds in the chart 1700 of Figure 17. Based on the determined load regions for UPF1 104 and UPF2 106, the weights determined by the SMF 102 for UPF1 104 load balancing are 1 & 4. For example, based on the detected predicted load of the new UE appearing on the network, if the SMF 102 determines that the predicted load of the new UE is 2 units of load, the SMF 102 will

perform load balancing such that there is a 10% probability the SMF 102 anchors the UE to UPF1 104 and a 90% probability the SMF 102 anchors the UE to UPF2 106. If the SMF 102 determines that the predicted load of the new UE is 1 unit of load, the SMF 102 will perform load balancing such that there is a 28.6% probability the SMF 102 anchors the UE to UPF1 104 and a 71.4% probability the SMF 102 anchors the UE to UPF2 106. Any new UE in the region of UPF2 106 (data center 2) are anchored on UPF2 106.

At time  $t=4$ , the SMF 102 determines that the current load of UPF1 104 is 75% of the total load capacity for UPF1 104 and thus determines the load region or UPF1 104 is "very high" based on the load thresholds in the chart 1700 of Figure 17. At time  $t=4$ , the SMF 102 also determines that the current load of UPF2 106 is 25% of the total load capacity for UPF2 106 and thus determines the load region or UPF2 106 is "low" based on the load thresholds in the chart 1700 of Figure 17. In this case, based on the determined load regions for UPF1 104 and UPF2 106, regardless of any predicted load of the new UE, any new UE in the region of UPF1 104 (data center 1) are anchored on UPF2 106 and any new UE in the region of UPF2 106 (data center 2) are also anchored on UPF2 106.

Figure 20 illustrates a chart 2000 showing an example of possible load-regions for non-low latency traffic that a current load of a UPF may be determined to fall within that may be used in the processes of Figures 8-10 in accordance with embodiments described herein.

The chart 2000 indicates the load region 2002 and the range of percentage of total load capacity 2004 for each load-region 2002 defined by respective load thresholds indicated in the chart 1700. The load-regions shown in the chart 2000 are used for UPF load balancing based on special considerations for low latency traffic. The load thresholds are reduced compared to those of chart 1700 in Figure 17 by a percentage amount of capacity dedicated for low-latency network traffic. In the present example, the percentage amount of capacity dedicated for low-latency network traffic is 10% and thus the load thresholds for non-low latency traffic are reduced by 10%. In particular, "Low" load region indicates a current UPF load of less than 20% of the total load capacity. The "Medium" load region indicates a current UPF load of greater than or equal to 20% and less than 45% of the total load capacity. The "High" load region indicates a current UPF load of greater than or equal to 45% and less than 65% of the total load capacity. The "Very High" load



region indicates a current UPF load of greater than or equal to 65% of the total load capacity. There may be additional or different load regions in various other embodiments. The chart 1700 may be stored or represented as a data structure in computer memory and may be maintained and/or accessible by the SMF 102.

5           In one embodiment, in performing UPF load balancing based on special considerations for low latency traffic, the SMF 102 may perform such load balancing in the manner described with respect to the processes of Figures 8-10 and the example herein described with respect to Figures 18, but uses the non-low latency load thresholds in chart 1700 instead that are reduced to dedicate 10% of the total  
10 load capacity for low-latency network traffic.

          Also, a chart and/or corresponding data structure similar to that of Figure 17 using load thresholds shown therein may be generated or used by the SMF 102 to perform UPF load balancing supporting multiple slices as described herein with respect to the processes of Figures 11-13. In such an embodiment, the chart and/or  
15 corresponding data structure also or instead indicates such load-thresholds for each slice depending on the UPF and network slice capacity. For example, these load thresholds may be percentage based, such as 30%, 55%, and 75% for slice 1 and 30%, 45%, and 65% for slice 2. Also, depending on the UPFs slice based load-regions, the SMF 102 creates slice based weights for a weighted load balancing  
20 between slices of different UPFs. The SMF 102 maintains the load-regions of the UPFs and creates the slice weights for the weighted load balancing of the UEs/PDU sessions among the UPFs. In such an embodiment, depending on the UPFs slice loads, the SMF 102 creates slice weights for a weighted load balancing of the UEs/PDU sessions among the UPFs.

25           In various embodiments, there may be additional UPFs and additional corresponding data centers. Additional load regions may be determined for each UPF, and the weighting and load balancing may be performed, in one example embodiment as described herein, to adjust the probabilities that a new UE appearing in the network is anchored on a particular UPF based on determined load regions for  
30 each UPF and the location of the UE. Such load balancing may instead be achieved using credit/token based scheduling (e.g., weighted round robin).

          Figure 21 shows a system diagram that describe various implementations of computing systems 2100 for implementing embodiments described herein.

The SMF 102 and the UPF, such as UPF1 104 and UPF2 106, can be implemented either as a network elements on dedicated hardware, as a software instance running on dedicated hardware, or as a virtualized function instantiated on an appropriate platform, e.g., a cloud infrastructure. In some embodiments, such

5 NFs may be completely software-based and designed as cloud-native, meaning that they're agnostic to the underlying cloud infrastructure, allowing higher deployment agility and flexibility. However, Figure 21 illustrates an example of underlying hardware on which the SMF 102 and the UPF, such as UPF1 104 and UPF2 106, may be implemented. For example, SMF 102 may be implemented using SMF

10 computer system(s) 2101. In some embodiments, one or more special-purpose computing systems may be used to implement SMF 102. Accordingly, various embodiments described herein may be implemented in software, hardware, firmware, or in some combination thereof. SMF computer system(s) 2101 may include memory 2102, one or more central processing units (CPUs) 2114, I/O

15 interfaces 2118, other computer-readable media 2120, and network connections 2122.

Memory 2102 may include one or more various types of non-volatile and/or volatile storage technologies. Examples of memory 2102 may include, but are not limited to, flash memory, hard disk drives, optical drives, solid-state drives, various

20 types of random access memory (RAM), various types of read-only memory (ROM), other computer-readable storage media (also referred to as processor-readable storage media), or the like, or any combination thereof. Memory 2102 may be utilized to store information, including computer-readable instructions that are utilized

25 Memory 2102 may have stored thereon SMF module 2104. The SMF module 2104 is configured to implement and/or perform some or all of the functions of the SMF 102 described herein. Memory 2102 may also store other programs and data 2110, which may include load thresholds, load-regions, databases, load-balancing rules, AI or ML programs to perform predictive analysis of UPF load based on

30 predicted UE throughput, CPU utilization and/or memory utilization using data from the NWDAF, user interfaces, operating systems, other network management functions, other NFs, etc.

Network connections 2122 are configured to communicate with other computing devices to facilitate the load balancing described herein. In various

embodiments, the network connections 2122 include transmitters and receivers (not illustrated) to send and receive data as described herein, such as sending data to and receiving data from UPFs, UEs and other NFs to send and receive instructions, commands and data to implement the processes described herein. I/O interfaces  
5 2118 may include a video interfaces, other data input or output interfaces, or the like. Other computer-readable media 2120 may include other types of stationary or removable computer-readable media, such as removable flash drives, external hard drives, or the like.

In some embodiments, one or more special-purpose computing systems may  
10 be used to implement UPF, such as UPF1 104 and UPF2 106. Accordingly, various embodiments described herein may be implemented in software, hardware, firmware, or in some combination thereof. UPF computer system(s) 2112 is an example of a computer system that may implement a UPF, such as UPF1 104 and UPF2 106. For example, computer system(s) 2112 may be present in data center 1  
15 to implement UPF1 104 or present in data center 2 to implement UPF2 106. Computer system(s) 2112 may include memory 2130, one or more central processing units (CPUs) 2144, I/O interfaces 2148, other computer-readable media 2150, and network connections 2152.

Memory 2130 may include one or more various types of non-volatile and/or  
20 volatile storage technologies similar to memory 2102. Memory 2130 may be utilized to store information, including computer-readable instructions that are utilized by CPU 2144 to perform actions, including embodiments described herein.

Memory 2130 may have stored thereon UPF module 2124. The UPF module  
25 214 receives the messages or instructions from the SMF module 204 to perform the load balancing operations as described herein. Memory 2130 may also store other programs and data 2138, which may include load thresholds, load-regions, databases, load-balancing rules, AI or ML programs to perform predictive analysis of UPF load based on predicted UE throughput, CPU utilization and/or memory utilization using data from the NWDAF, user interfaces, operating systems, other  
30 network management functions, other NFs, etc.

Network connections 2152 are configured to communicate with other computing devices, such as SMF computer system(s) 2101. In various embodiments, the network connections 2152 include transmitters and receivers (not illustrated) to send and receive data as described herein. I/O interfaces 2148 may

include one or more other data input or output interfaces. Other computer-readable media 2150 may include other types of stationary or removable computer-readable media, such as removable flash drives, external hard drives, or the like.

The various embodiments described above can be combined to provide  
5 further embodiments. These and other changes can be made to the embodiments in  
light of the above-detailed description. In general, in the following claims, the terms  
used should not be construed to limit the claims to the specific embodiments  
disclosed in the specification and the claims, but should be construed to include all  
possible embodiments along with the full scope of equivalents to which such claims  
10 are entitled. Accordingly, the claims are not limited by the disclosure.

## CLAIMS

1. A system, comprising:
  - a memory that stores computer instructions; and
  - a processor that executes the computer instructions to perform actions, the actions including:
    - maintaining a plurality load thresholds for each user plane function (UPF) of a plurality of UPFs in a cellular telecommunication network, wherein:
      - the plurality of UPFs serve as anchor points between user equipment (UE) in the cellular telecommunication network and a data network (DN);
      - each UPF of the plurality of UPFs is a virtual network function responsible for interconnecting packet data unit (PDU) sessions between the user equipment (UE) and the DN by anchoring the PDU sessions on individual UPFs;
      - the plurality of load thresholds for each UPF depend on a respective capacity of each UPF to have PDU sessions anchored thereon; and
      - an amount of load put on a UPF by a UE appearing in the cellular telecommunication network is assumed to be identical for all UEs appearing in the cellular telecommunication network;
      - receiving a request to anchor on a UPF a PDU session of a new UE newly appearing on the cellular telecommunication network;
      - selecting a UPF of the plurality of UPFs on which to anchor the PDU session based on a location of the new UE and determined load-regions for each UPF of the plurality of UPFs defined by the plurality of load thresholds; and
      - anchoring the PDU session of the new UE to the selected UPF.
2. The system of claim 1, wherein the selecting a UPF of the plurality of UPFs includes:
  - generating weights for selecting the UPF based on the determined load-regions; and
  - selecting the UPF based on the generated weights.
3. The system of claim 2, wherein the selecting a UPF of the plurality of UPFs includes:

generating multiple load-regions, wherein each load-region corresponds to a different range of current load of a UPF defined by one or more of lower and upper threshold percentages of UPF load capacity;

in response to receiving the request to anchor the PDU session, for each UPF of the plurality of UPFs, determining a load region from the multiple load-regions that a current load of the UPF falls within; and

selecting a UPF of the plurality of UPFs based on the determined load-regions for the plurality of UPFs and the weights generated based on the determined load regions.

4. The system of claim 3 wherein each UPF of the plurality of UPFs is associated with a different respective geographic area and selecting a UPF of the plurality of UPFs based on the generated weights and the determined load-regions for the plurality of UPFs includes:

determining a particular UPF of the plurality of UPFs has a respective geographic area within which the location of the new UE falls; and

selecting the particular UPF in response to the determined load-region of the particular UPF being a load-region indicating a current load of the particular UPF is below a threshold capacity.

5. The system of claim 3, wherein each UPF of the plurality of UPFs is associated with a different respective geographic area and the selecting a UPF of the plurality of UPFs based on the generated weights and the determined load-regions for the plurality of UPFs includes:

determining a particular UPF of the plurality of UPFs has a respective geographic area within which the location of the new UE falls;

determining whether the particular UPF has a determined load region indicating a current load of the particular UPF is in a different load region indicating a higher current load of the particular UPF than a current load of another UPF of the plurality of UPFs; and

in response to the particular UPF having a determined load region indicating the current load of the particular UPF is in a different load region indicating a higher current load of the particular UPF than the current load of another UPF of the

plurality of UPFs, weighting the selection of a UPF of the plurality of UPFs such that the particular UPF is not overloaded compared to the other UPF of the plurality of UPFs by using credit/token-based weighted scheduling or probability-based weighted scheduling.

6. The system of claim 3, wherein:

each UPF of the plurality of UPFs is associated with a different respective geographic area;

the selecting a UPF of the plurality of UPFs based on the generated weights and the determined load-regions for the plurality of UPFs includes determining a particular UPF of the plurality of UPFs has a respective geographic area within which the location of the new UE falls; and

the selection of the UPF is weighted by using credit/token-based weighted scheduling or probability-based weighted scheduling such that the frequency of selection of the particular UPF decreases as a difference between a higher current load of the particular UPF and a lower current load of at least one UPF of the plurality of UPFs increases, as indicated by the load regions determined for each UPF of the plurality of UPFs.

7. The system of claim 3 wherein the generating multiple load-regions includes:

generating a lowest load-region indicating a current UPF load less than a first threshold percentage of UPF capacity;

generating one or more intermediate non-overlapping load-regions each defined by respective lower and upper threshold percentages of UPF capacity and indicating a current load greater than the lowest load-region; and

generating a highest load-region indicating a current UPF load greater than a second threshold percentage of UPF capacity and greater than the one or more intermediate non-overlapping load-regions.

8. A method, comprising:

electronically maintaining a plurality of load thresholds for each user plane function (UPF) of a plurality of UPFs in a cellular telecommunication network, wherein:

the plurality of UPFs serve as anchor points between user equipment (UE) in the cellular telecommunication network and a data network (DN);

each UPF of the plurality of UPFs is a virtual network function responsible for interconnecting packet data unit (PDU) sessions between the user equipment (UE) and the DN by anchoring the PDU sessions on individual UPFs;

the plurality of load thresholds for each UPF depend on a respective capacity of each UPF to have PDU sessions anchored thereon; and

an amount of load put on a UPF by a UE appearing in the cellular telecommunication network is assumed to be identical for all UEs appearing in the cellular telecommunication network;

electronically receiving a request to anchor on a UPF a PDU session of a new UE newly appearing on the cellular telecommunication network;

electronically selecting a UPF of the plurality of UPFs on which to anchor the PDU session based on a location of the new UE and determined load-regions for each UPF of the plurality of UPFs defined by the plurality of load thresholds; and

electronically anchoring the PDU session of the new UE to the selected UPF.

9. The method of claim 8, wherein the selecting a UPF of the plurality of UPFs includes:

generating weights for selecting the UPF based on the determined load-regions; and

selecting the UPF based on the generated weights.

10. The method of claim 9, wherein the selecting a UPF of the plurality of UPFs includes:

generating multiple load-regions, wherein each load-region corresponds to a different range of current load of a UPF defined by one or more of lower and upper threshold percentages of UPF load capacity;



in response to receiving the request to anchor the PDU session, for each UPF of the plurality of UPFs, determining a load region from the multiple load-regions that a current load of the UPF falls within; and

selecting a UPF of the plurality of UPFs based on the determined load-regions for the plurality of UPFs and the weights generated based on the determined load regions.

11. The method of claim 10 wherein each UPF of the plurality of UPFs is associated with a different respective geographic area and selecting a UPF of the plurality of UPFs based on the generated weights and the determined load-regions for the plurality of UPFs includes:

determining a particular UPF of the plurality of UPFs has a respective geographic area within which the location of the new UE falls; and

selecting the particular UPF in response to the determined load-region of the particular UPF being a load-region indicating a current load of the particular UPF is below a threshold capacity.

12. The method of claim 10, wherein each UPF of the plurality of UPFs is associated with a different respective geographic area and the selecting a UPF of the plurality of UPFs based on the generated weights and the determined load-regions for the plurality of UPFs includes:

determining a particular UPF of the plurality of UPFs has a respective geographic area within which the location of the new UE falls;

determining whether the particular UPF has a determined load region indicating a current load of the particular UPF is in a different load region indicating a higher current load of the particular UPF than a current load of another UPF of the plurality of UPFs; and

in response to the particular UPF having a determined load region indicating the current load of the particular UPF is in a different load region indicating a higher current load of the particular UPF than the current load of another UPF of the plurality of UPFs, weighting the selection of a UPF of the plurality of UPFs such that the particular UPF is not overloaded compared to the other UPF of the plurality of

UPFs by using credit/token-based weighted scheduling or probability-based weighted scheduling.

13. The method of claim 10, wherein:

each UPF of the plurality of UPFs is associated with a different respective geographic area;

the selecting a UPF of the plurality of UPFs based on the generated weights and the determined load-regions for the plurality of UPFs includes determining a particular UPF of the plurality of UPFs has a respective geographic area within which the location of the new UE falls; and

the selection of the UPF is weighted by using credit/token-based weighted scheduling or probability-based weighted scheduling such that the frequency of selection of the particular UPF decreases as a difference between a higher current load of the particular UPF and a lower current load of at least one UPF of the plurality of UPFs increases, as indicated by the load regions determined for each UPF of the plurality of UPFs.

14. The method of claim 10 wherein the generating multiple load-regions includes:

generating a lowest load-region indicating a current UPF load less than a first threshold percentage of UPF capacity;

generating one or more intermediate non-overlapping load-regions each defined by respective lower and upper threshold percentages of UPF capacity and indicating a current load greater than the lowest load-region; and

generating a highest load-region indicating a current UPF load greater than a second threshold percentage of UPF capacity and greater than the one or more intermediate non-overlapping load-regions.

15. A non-transitory computer-readable storage medium having computer-executable instructions stored thereon that, when executed by at least one computer processor, cause actions to be performed including:

maintaining a plurality of load thresholds for each user plane function (UPF) of a plurality of UPFs in a cellular telecommunication network, wherein:

the plurality of UPFs serve as anchor points between user equipment (UE) in the cellular telecommunication network and a data network (DN);

each UPF of the plurality of UPFs is a virtual network function responsible for interconnecting packet data unit (PDU) sessions between the user equipment (UE) and the DN by anchoring the PDU sessions on individual UPFs;

the plurality of load thresholds for each UPF depend on a respective capacity of each UPF to have PDU sessions anchored thereon; and

an amount of load put on a UPF by a UE appearing in the cellular telecommunication network is assumed to be identical for all UEs appearing in the cellular telecommunication network;

receiving a request to anchor on a UPF a PDU session of a new UE newly appearing on the cellular telecommunication network;

selecting a UPF of the plurality of UPFs on which to anchor the PDU session based on a location of the new UE and determined load-regions for each UPF of the plurality of UPFs defined by the plurality of load thresholds; and

anchoring the PDU session of the new UE to the selected UPF.

16. The non-transitory computer-readable storage medium of claim 15, wherein the selecting a UPF of the plurality of UPFs includes:

generating weights for selecting the UPF based on the determined load-regions; and

selecting the UPF based on the generated weights.

17. The non-transitory computer-readable storage medium of claim 16, wherein the selecting a UPF of the plurality of UPFs includes:

generating multiple load-regions, wherein each load-region corresponds to a different range of current load of a UPF defined by one or more of lower and upper threshold percentages of UPF load capacity;

in response to receiving the request to anchor the PDU session, for each UPF of the plurality of UPFs, determining a load region from the multiple load-regions that a current load of the UPF falls within; and

selecting a UPF of the plurality of UPFs based on the determined load-regions for the plurality of UPFs and the weights generated based on the determined load regions.

18. The non-transitory computer-readable storage medium of claim 17 wherein each UPF of the plurality of UPFs is associated with a different respective geographic area and selecting a UPF of the plurality of UPFs based on the generated weights and the determined load-regions for the plurality of UPFs includes:

determining a particular UPF of the plurality of UPFs has a respective geographic area within which the location of the new UE falls; and

selecting the particular UPF in response to the determined load-region of the particular UPF being a load-region indicating a current load of the particular UPF is below a threshold capacity.

19. The non-transitory computer-readable storage medium of claim 17, wherein each UPF of the plurality of UPFs is associated with a different respective geographic area and the selecting a UPF of the plurality of UPFs based on the generated weights and the determined load-regions for the plurality of UPFs includes:

determining a particular UPF of the plurality of UPFs has a respective geographic area within which the location of the new UE falls;

determining whether the particular UPF has a determined load region indicating a current load of the particular UPF is in a different load region indicating a higher current load of the particular UPF than a current load of another UPF of the plurality of UPFs; and

in response to the particular UPF having a determined load region indicating the current load of the particular UPF is in a different load region indicating a higher current load of the particular UPF than the current load of another UPF of the plurality of UPFs, weighting the selection of a UPF of the plurality of UPFs such that the particular UPF is not overloaded compared to the other UPF of the plurality of UPFs by using credit/token-based weighted scheduling or probability-based weighted scheduling.

20. The non-transitory computer-readable storage medium of claim 17, wherein:

each UPF of the plurality of UPFs is associated with a different respective geographic area;

the selecting a UPF of the plurality of UPFs based on the generated weights and the determined load-regions for the plurality of UPFs includes determining a particular UPF of the plurality of UPFs has a respective geographic area within which the location of the new UE falls; and

the selection of the UPF is weighted by using credit/token-based weighted scheduling or probability-based weighted scheduling such that the frequency of selection of the particular UPF decreases as a difference between a higher current load of the particular UPF and a lower current load of at least one UPF of the plurality of UPFs increases, as indicated by the load regions determined for each UPF of the plurality of UPFs.


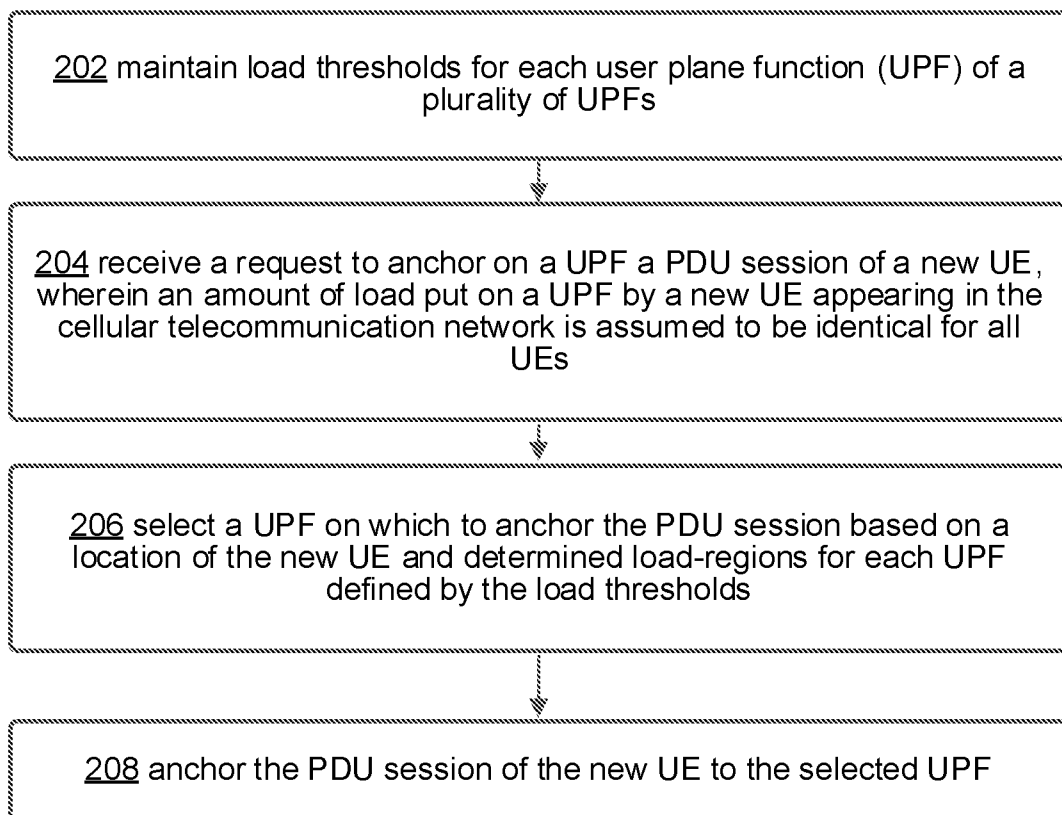


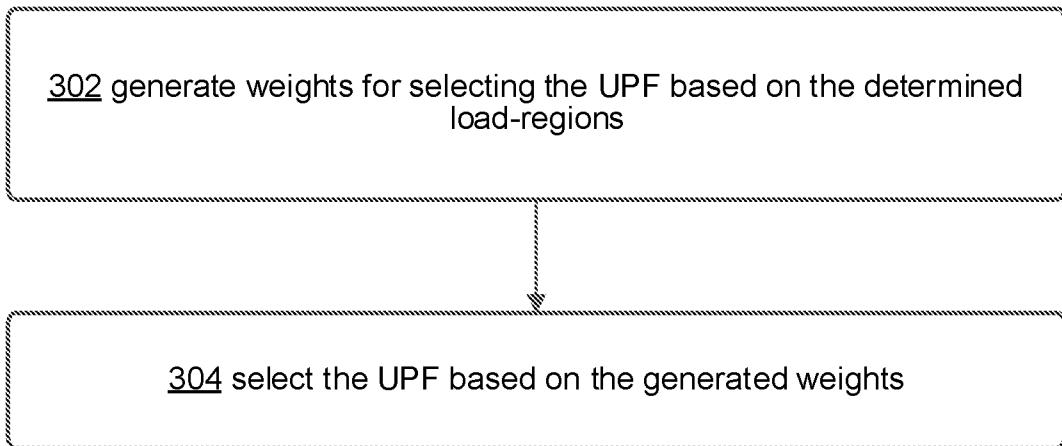
FIG. 1

200



**FIG. 2**

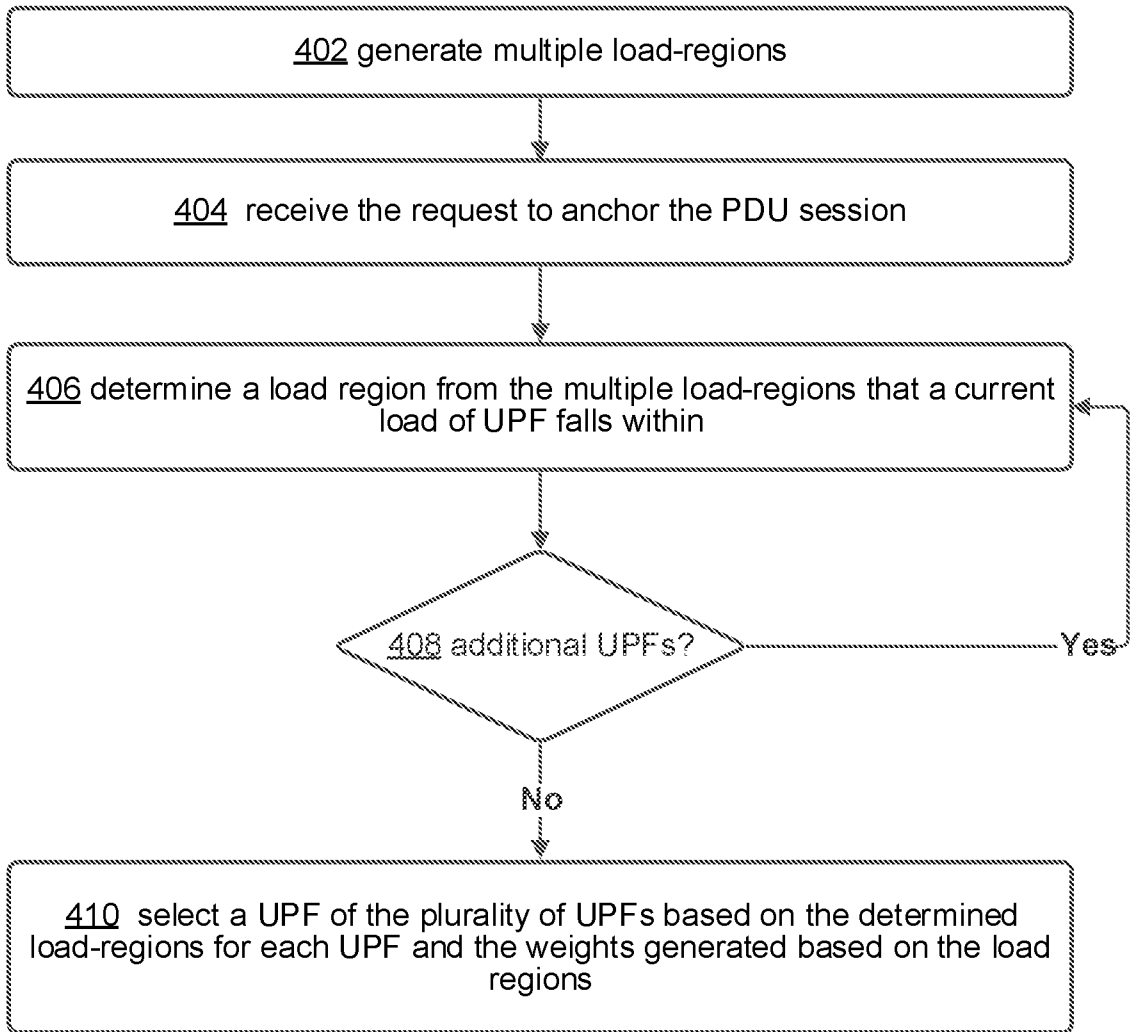
300



**FIG. 3**

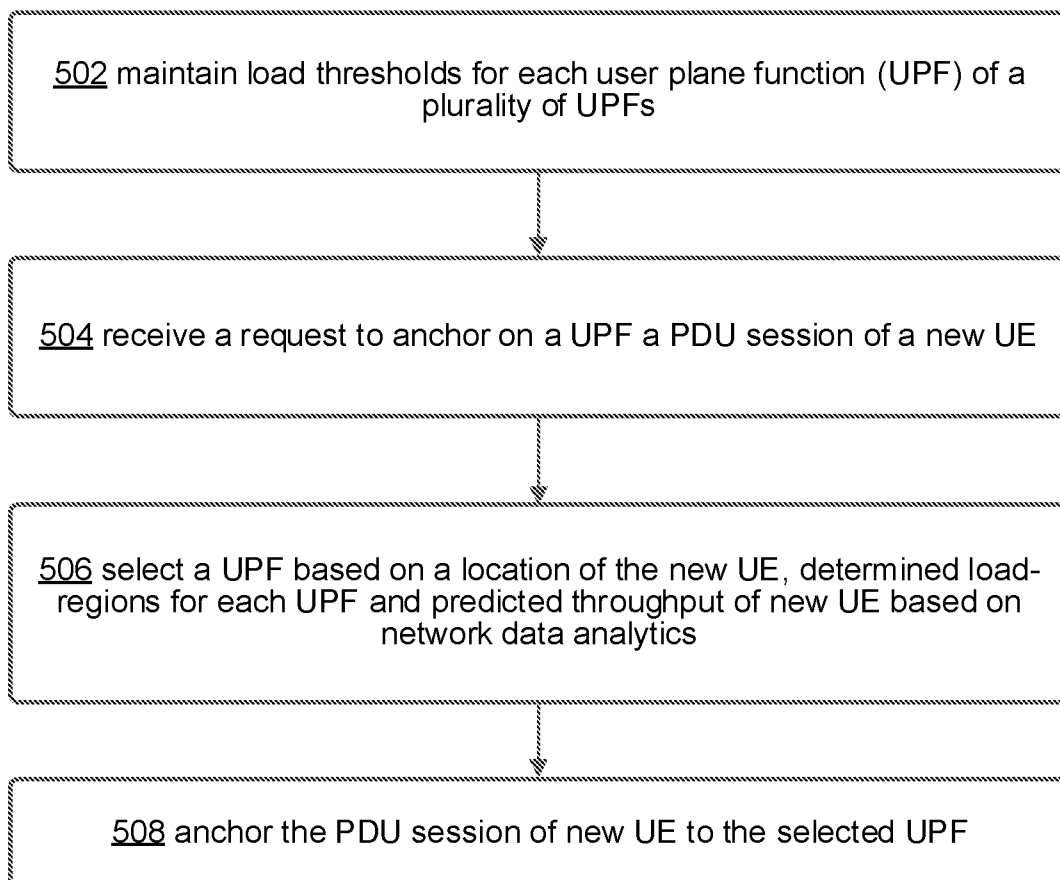


400

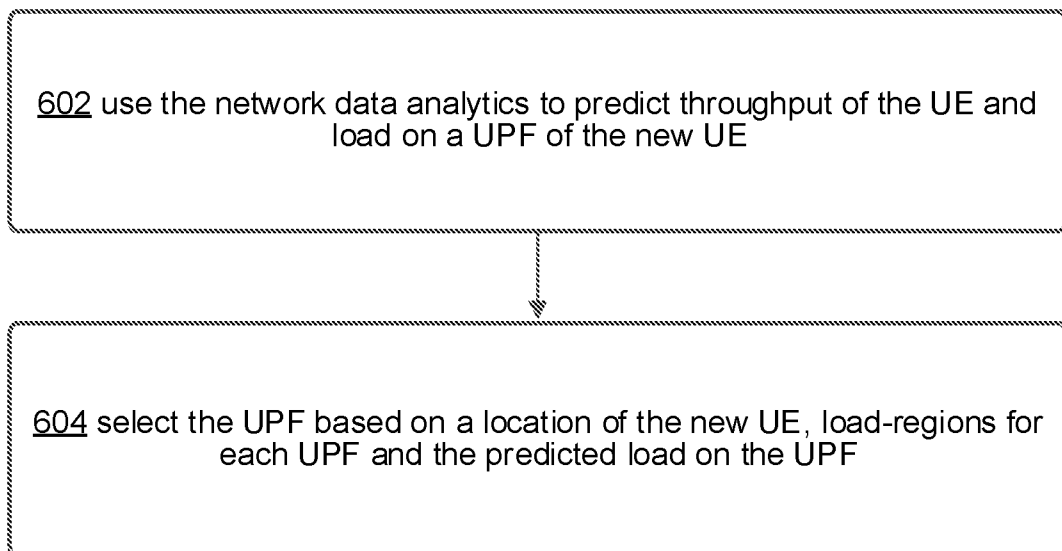


**FIG. 4**

500

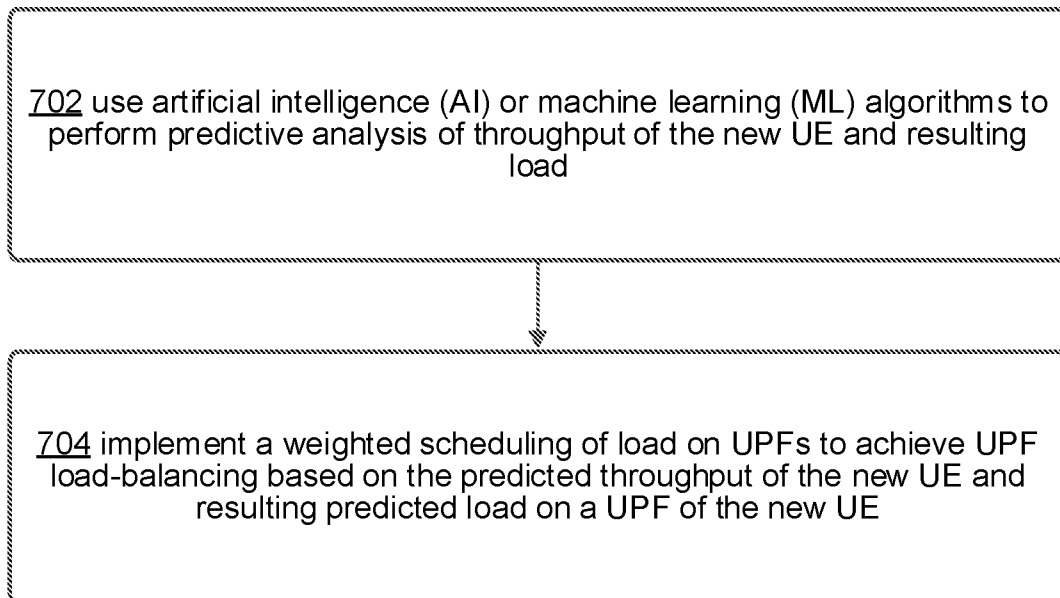
**FIG. 5**

600



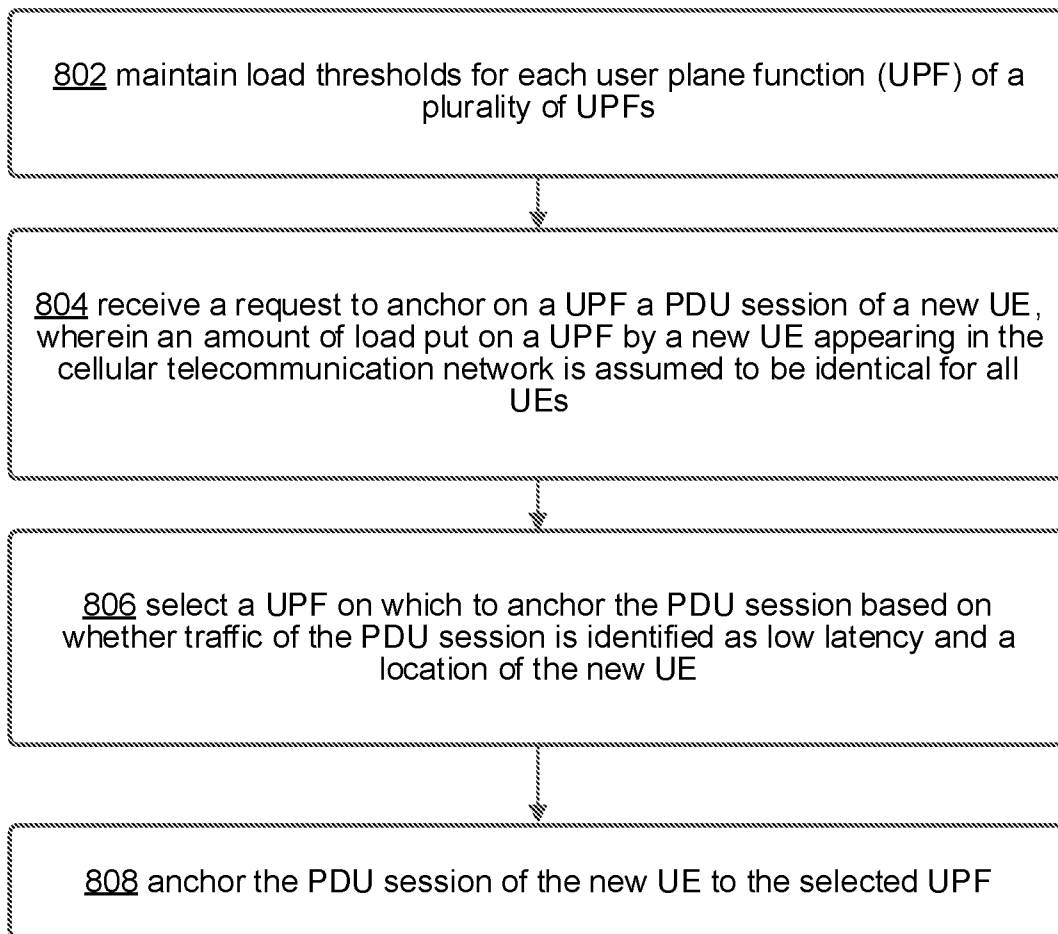
**FIG. 6**

700

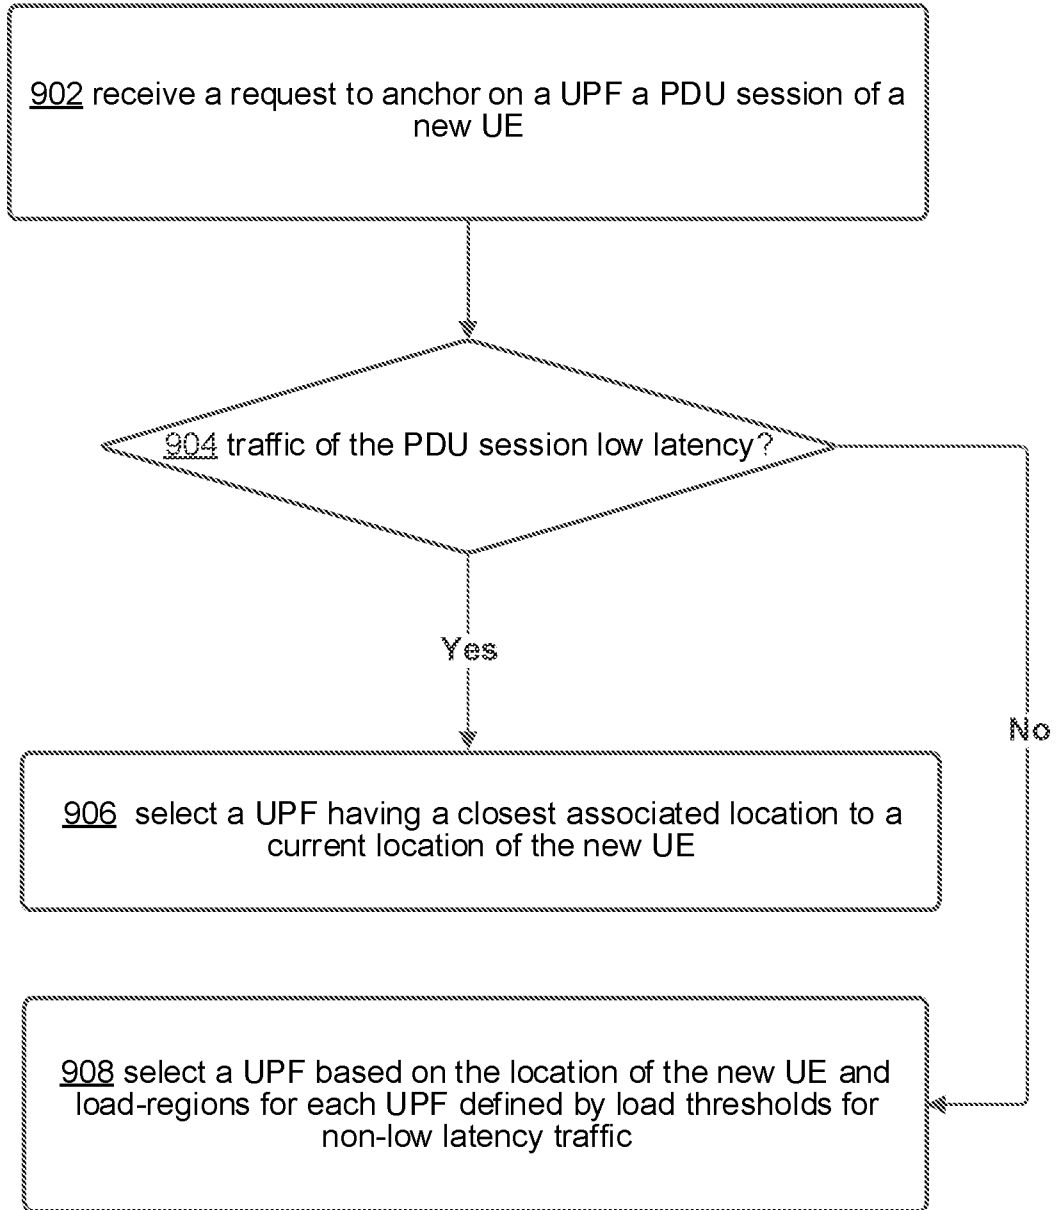


**FIG. 7**

700

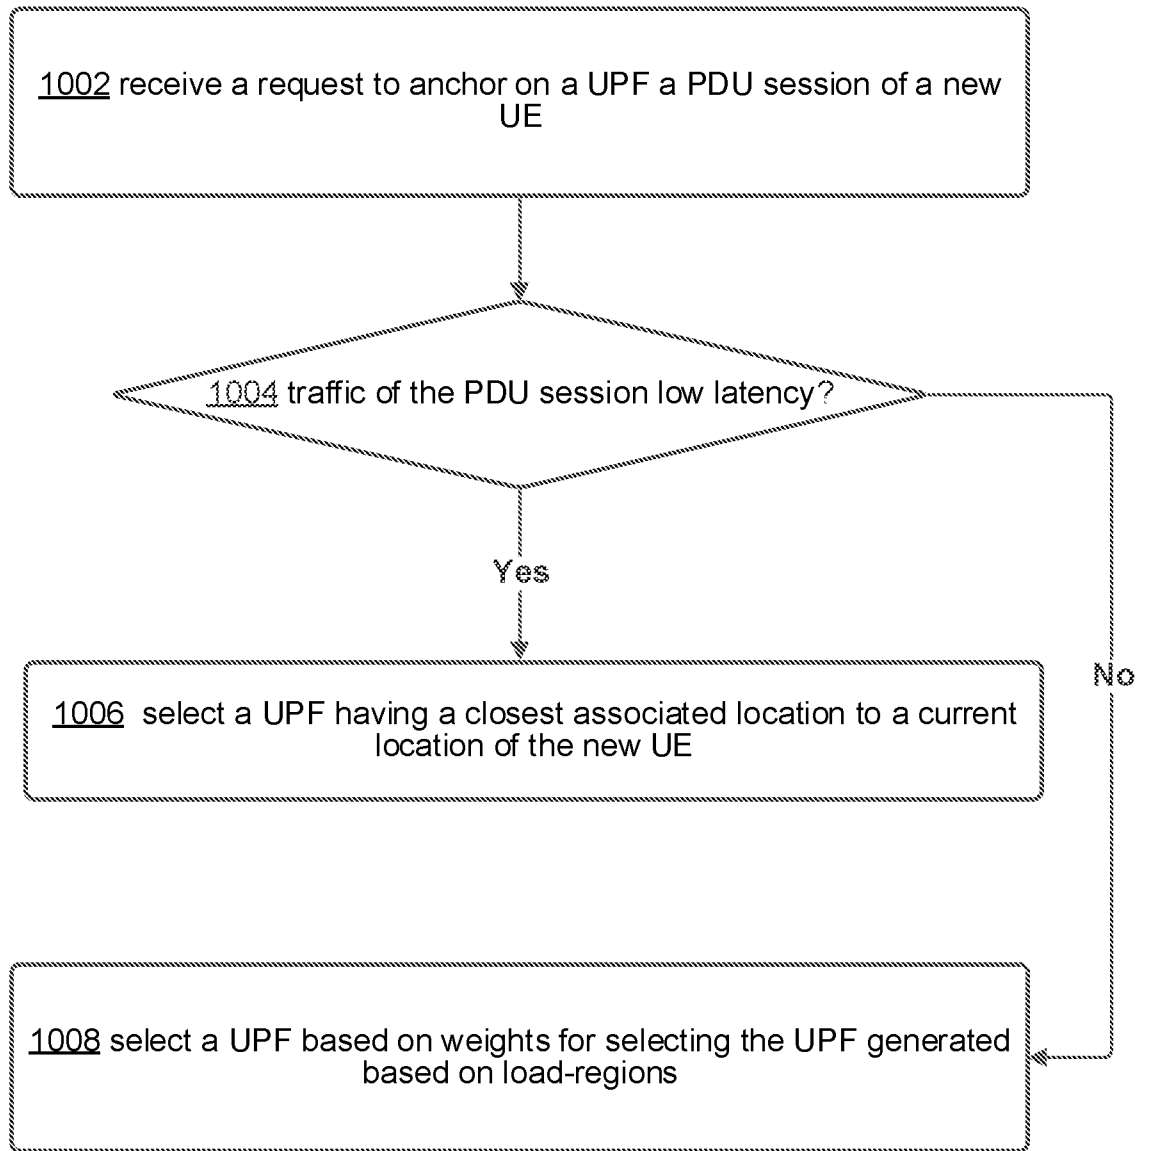
**FIG. 8**

900



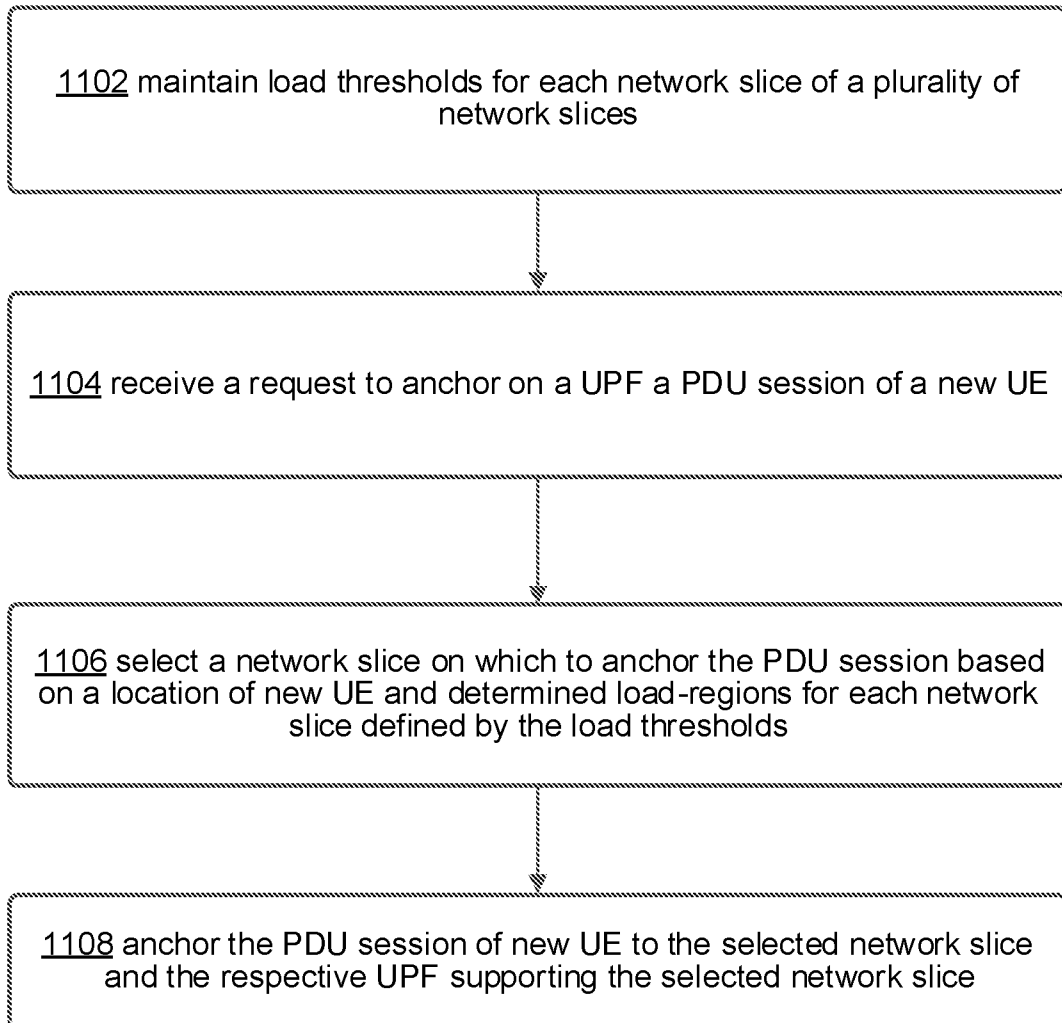
**FIG. 9**

1000



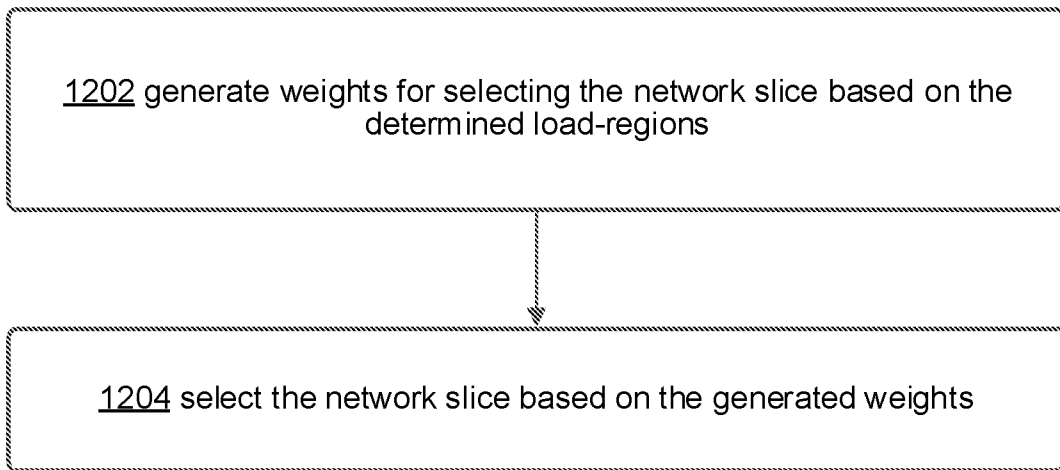

**FIG. 10**

1100

**FIG. 11**

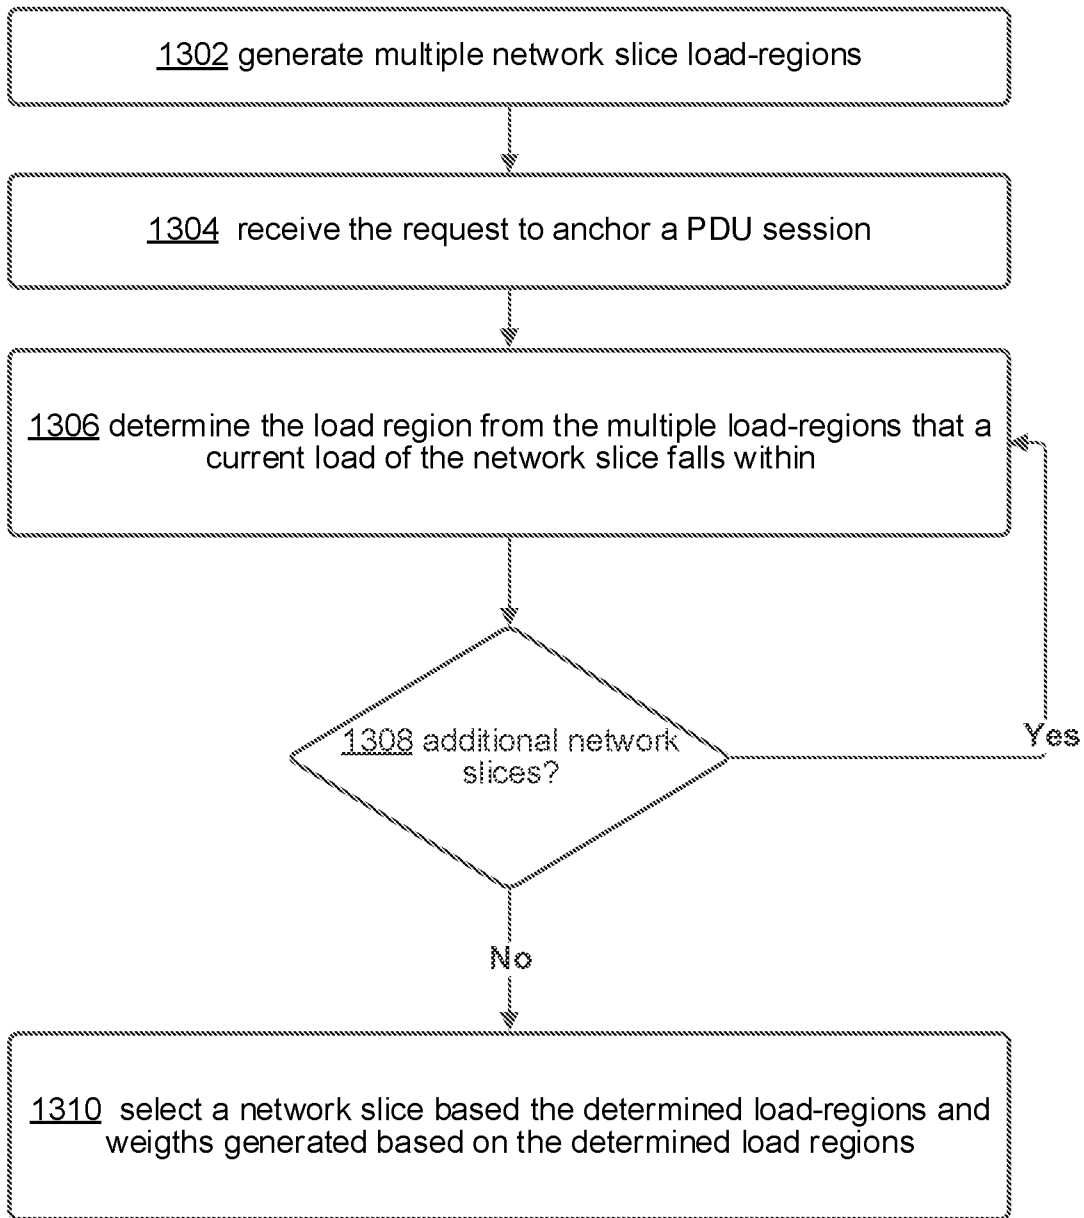


1200

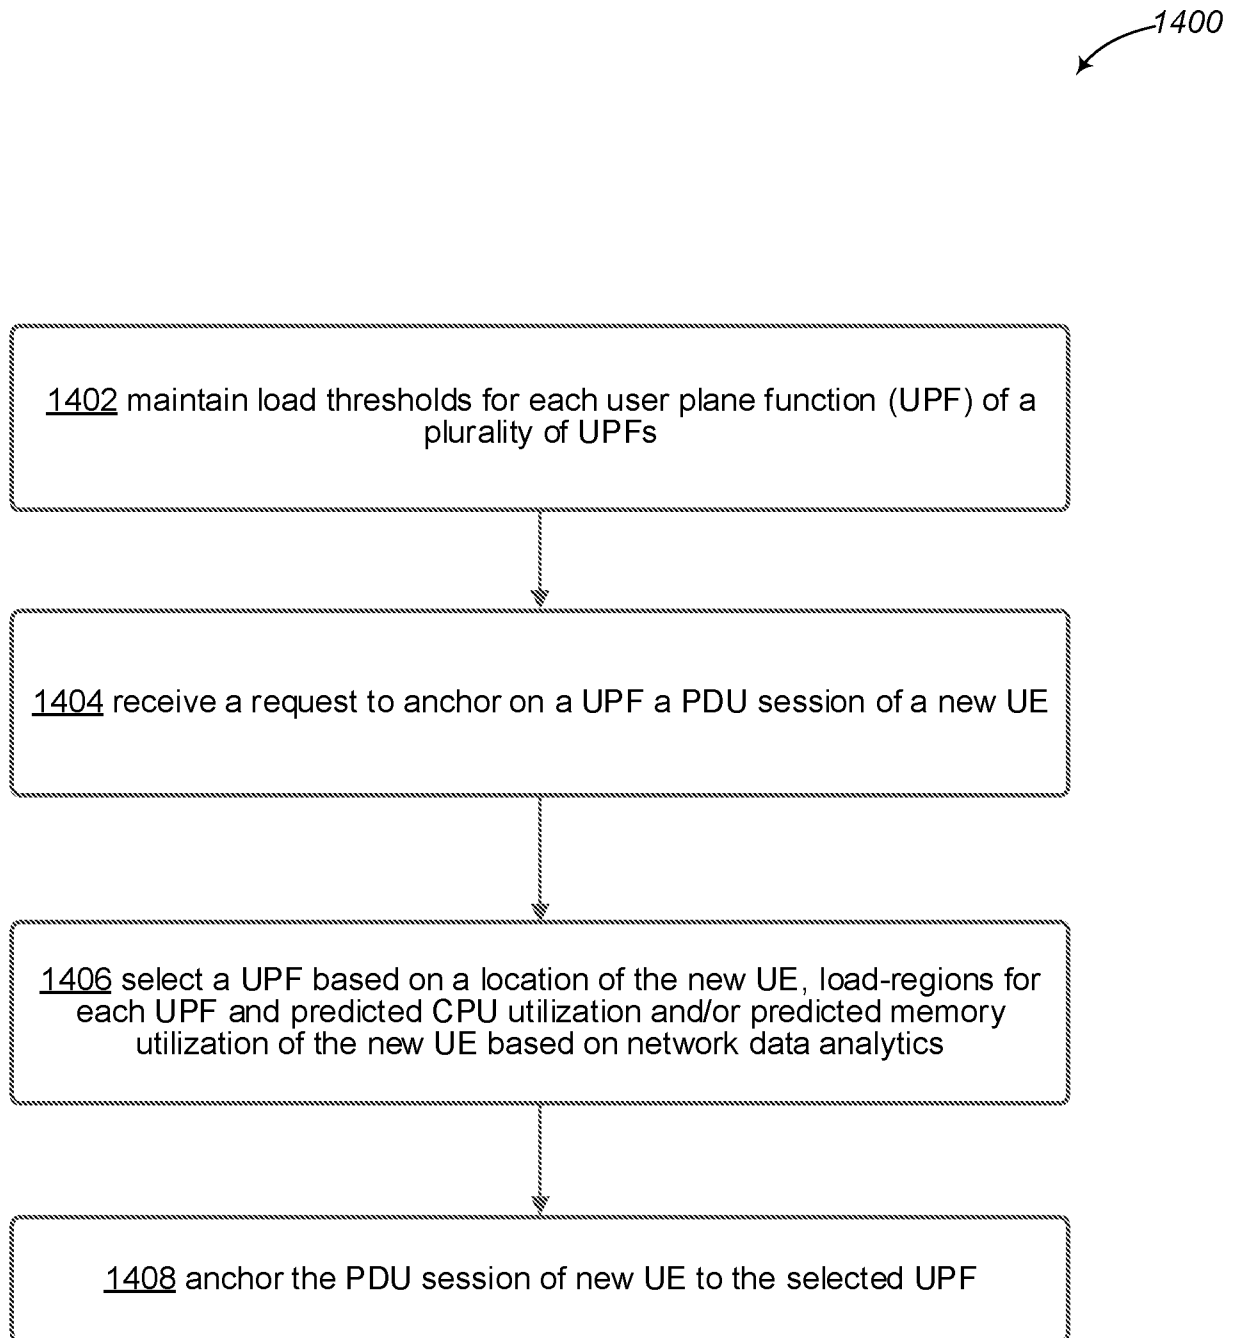


**FIG. 12**

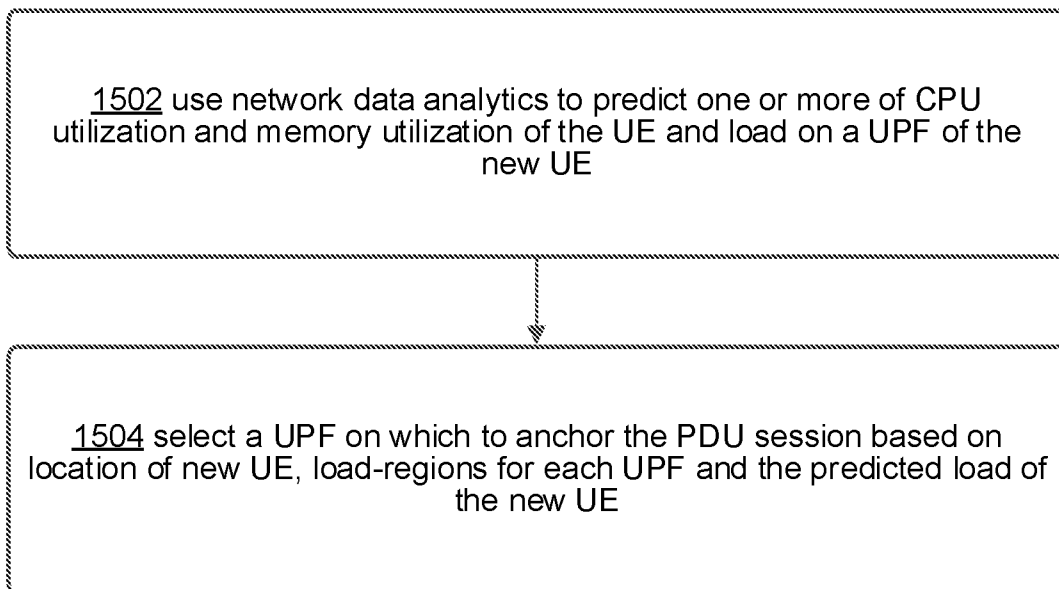
1300



**FIG. 13**

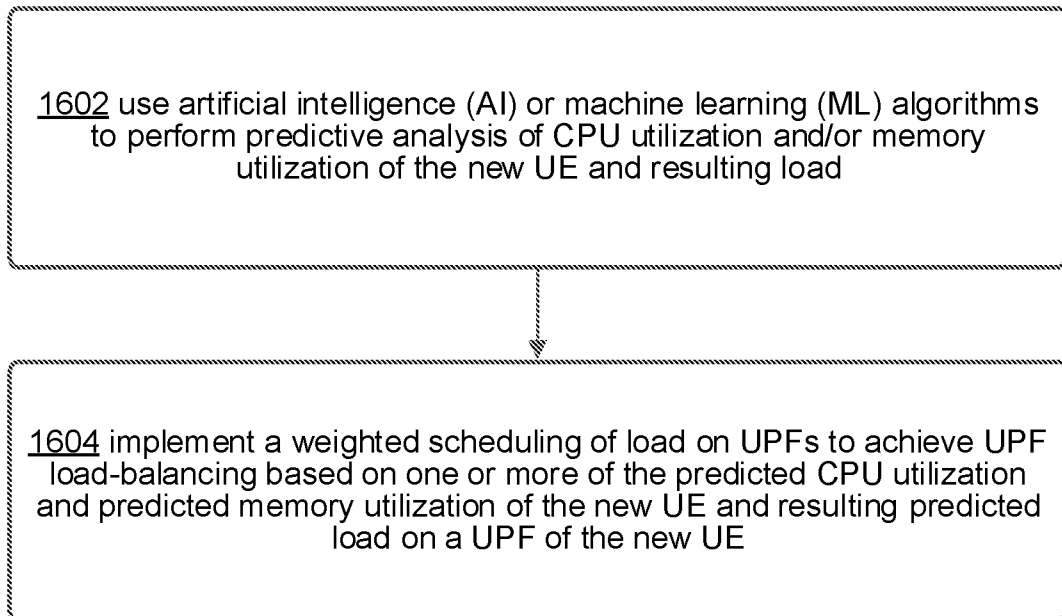
**FIG. 14**

1500



**FIG. 15**

1600



**FIG. 16**

1700

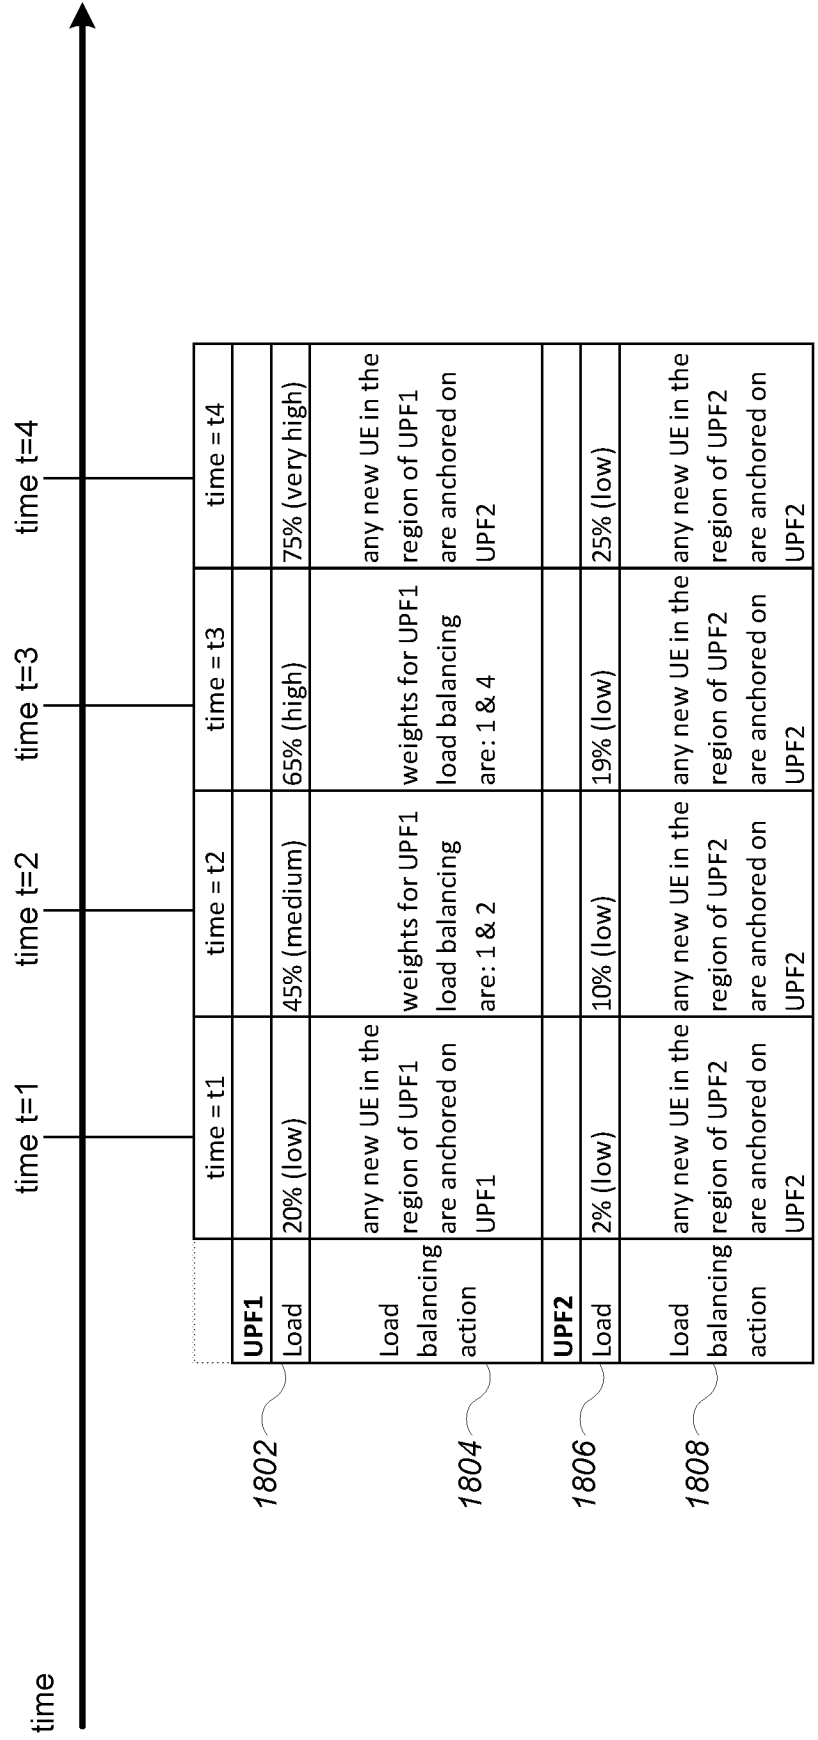
Load region	Percentage of Total Capacity
Low	< 30%
Medium	≥ 30% and < 55%
High	≥ 55% and < 75%
Very High	≥ 75%

1702

1704

FIG. 17

1800



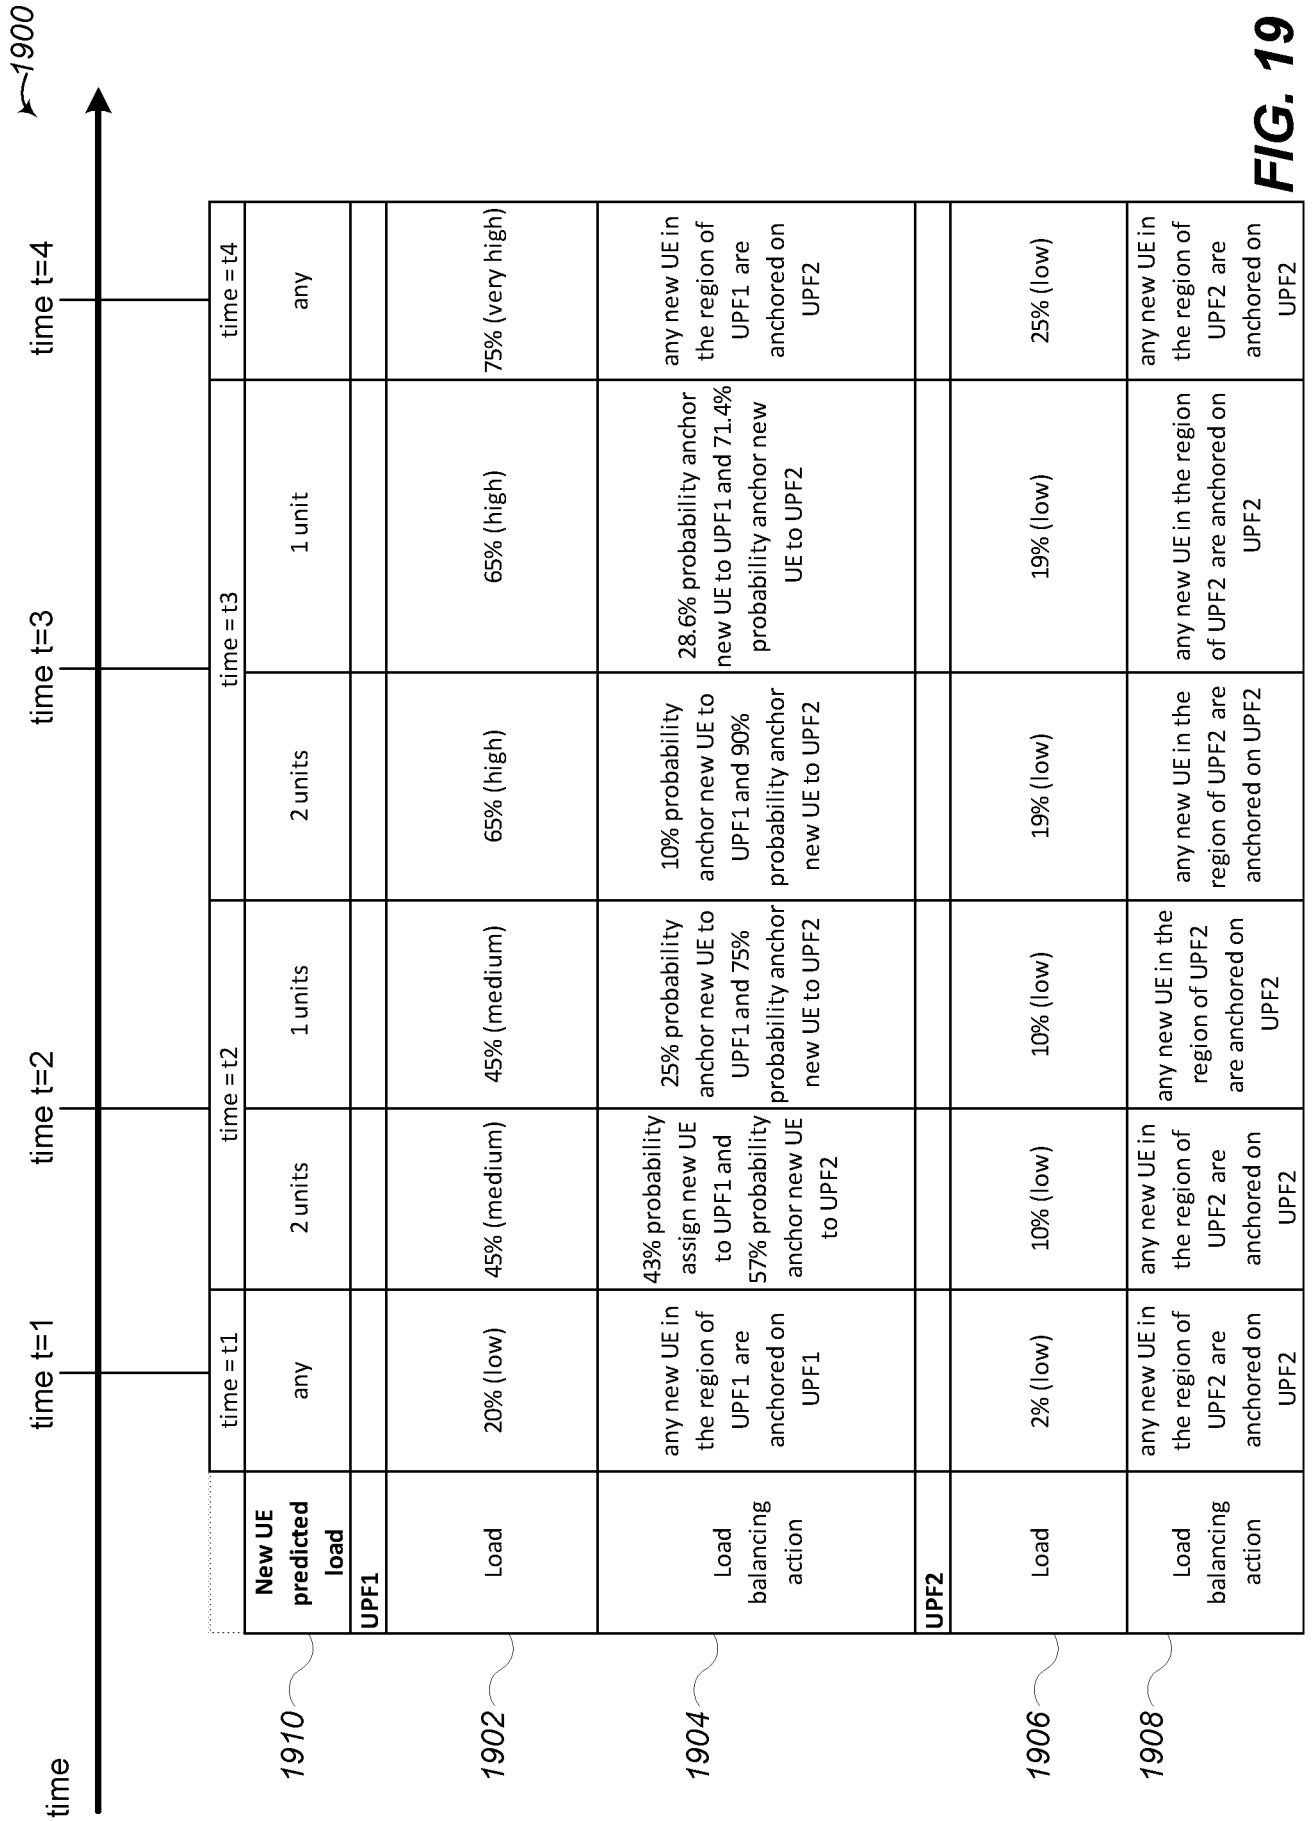
1802

1804

1806

1808

**FIG. 18**



**FIG. 19**



← 2000

# Non-low Latency Traffic Load-Regions\*

Load region	Percentage of Total Capacity
Low	< 20%
Medium	≥ 20% and < 45%
High	≥ 45% and < 65%
Very High	≥ 65%

\*10% of UPF capacity dedicated for low latency PDU sessions

**FIG. 20**

2100

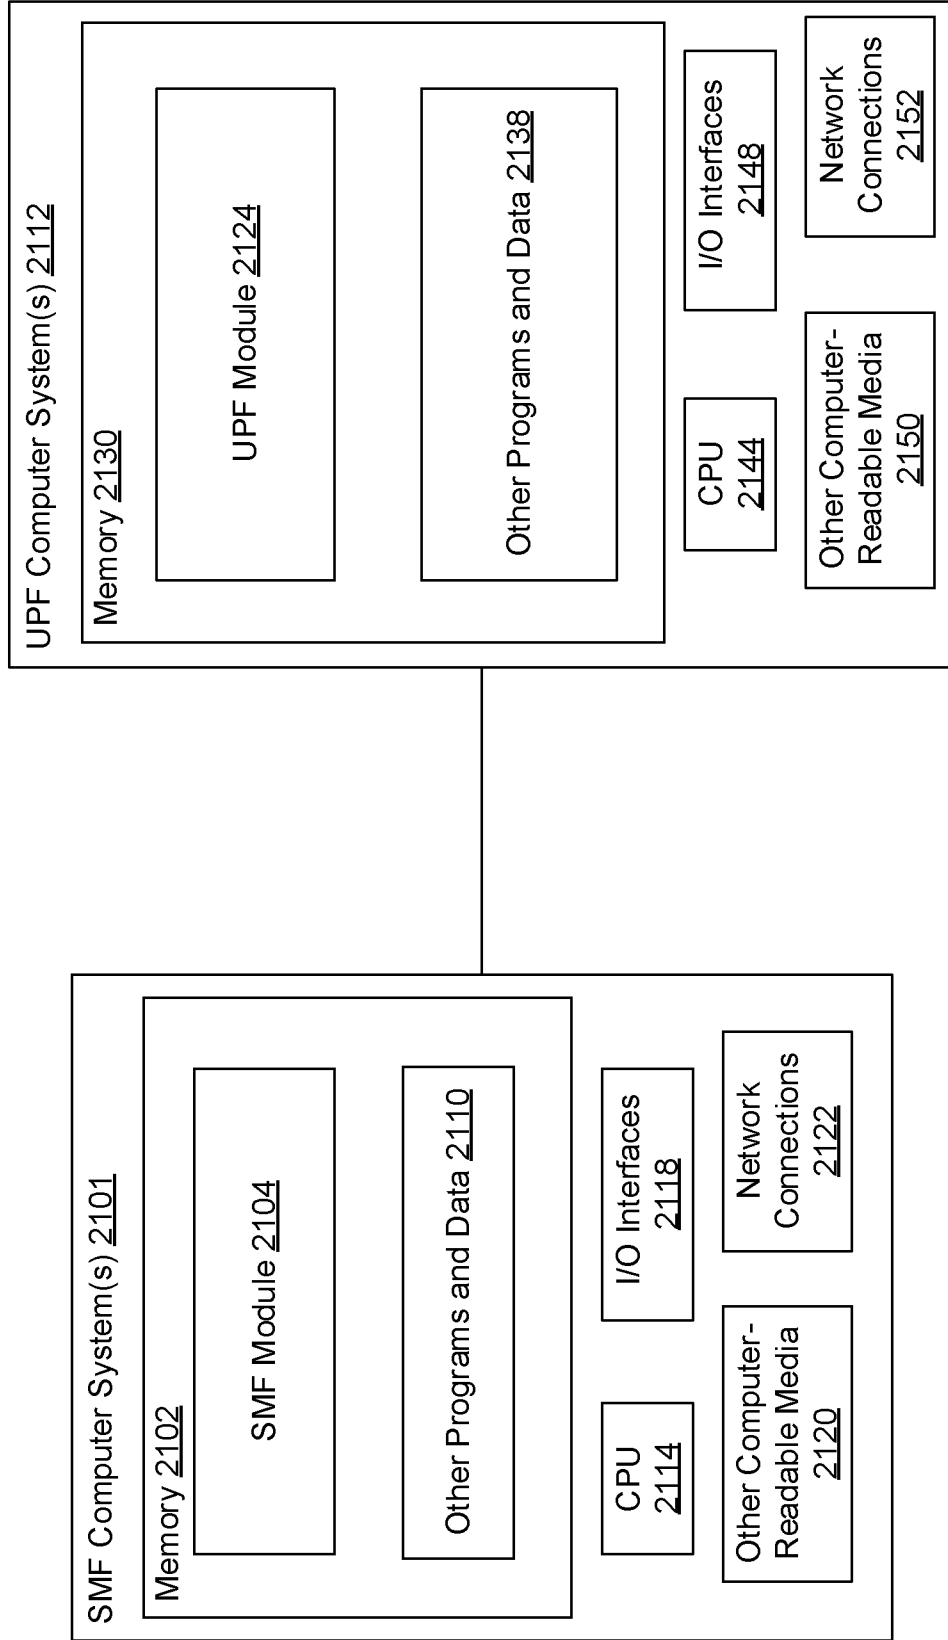


FIG. 21



## INTERNATIONAL SEARCH REPORT

International application No

PCT/US2022/040157

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	<p>"3rd Generation Partnership Project; Technical Specification Group Services and System Aspects; System architecture for the 5G System (5GS); Stage 2 (Release 16)", 3GPP STANDARD; 3GPP TS 23.501, 3RD GENERATION PARTNERSHIP PROJECT (3GPP), MOBILE COMPETENCE CENTRE ; 650, ROUTE DES LUCIOLES ; F-06921 SOPHIA-ANTIPOLIS CEDEX ; FRANCE , vol. SA WG2, no. V16.9.0 24 June 2021 (2021-06-24), pages 1-452, XP052029598, Retrieved from the Internet: URL:<a href="https://ftp.3gpp.org/Specs/archive/23_series/23.501/23501-g90.zip">https://ftp.3gpp.org/Specs/archive/23_ series/23.501/23501-g90.zip</a> [retrieved on 2021-06-24] chapter 6.3.3.3 -----</p>	1-20

# INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

**PCT/US2022/040157**

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
<b>CN 110662260</b>	<b>A</b>	<b>07-01-2020</b>	<b>NONE</b>
-----			
<b>CN 112867050</b>	<b>A</b>	<b>28-05-2021</b>	<b>NONE</b>
-----			