Mesh Networks

Technical Field

[0001] The present invention relates to mesh networks including mobile base stations, and methods of operation.

Background Art

[0002] Mesh networks may include mobile base stations, located on mobile platforms such as maritime vessels, land vehicles or aircraft for example. The mobile base stations may provide wireless (e.g. cellular) access to the network for terminals located on the mobile platform, or off the platform but within range of the mobile base station.

[0003] One problem with such mesh networks is that the position and/or orientation of the mobile base stations may not be under the control of the network operator. This may lead to regulatory or other compliance problems, for example relating to the location and/or spectrum usage of radio frequency (RF) transmissions from the mobile base stations.

Statements of the Invention

[0004] Aspects of the invention are defined in the accompanying claims.

[0005] In an embodiment, a method of controlling emissions from a mobile node in a mesh network comprises obtaining positional data relating to the mobile node, comparing the positional data with a geographical emissions restriction requirement, determining a beam mask for the mobile node so as to meet the geographical emissions restriction requirement and controlling the emissions of the mobile node based on the beam mask. The method may be performed by a controller remote from the mobile node. The controller may have out of band connectivity with the mobile node via one or more alternative networks other than the mesh network.

Brief Description of the Drawings

[0006] Specific embodiments of the present invention will now be described with reference to the accompanying drawings listed below.

Fig. 1 is diagram illustrating a scenario in which transmissions from a mobile base station may infringe a regulatory border.

Fig. 2 is a diagram of a network architecture in which embodiments of the invention may be implemented.

Fig. 3 is a diagram illustrating a blockage zone within antenna coverage of a mobile platform.

Fig. 4 shows an example of an antenna beam pattern of a mobile platform.

Fig. 5 is a diagram of a network architecture using multiple networks, in which embodiments of the invention may be implemented.

Fig. 6 is a flow diagram showing a method of beam masking in a first embodiment.

Fig. 7 is a flow diagram of a method of beam masking in a second embodiment.

Fig. 8 is a diagram illustrating a use case of an embodiment of the invention.

Fig. 9 is a diagram showing a first example of calculation of a beam mask in an embodiment.

Fig. 10 is a diagram showing a second example of calculation of a beam mask in an embodiment.

Fig. 11 is a diagram of a computer system on which functionalities of embodiments may be performed.

Detailed Description of Embodiments

[0007] A specific embodiment relates to a mesh network, such as but not limited to a 5G mesh network, including a plurality of interconnected nodes. The nodes may comprise mobile base stations which may provide connectivity to terminals and/or to other base stations.

[0008] In the scenario shown in Fig. 1, a mesh network 10 comprises a plurality of mobile base stations 1 located on respective mobile platforms 2. The mobile platforms 2 may comprise one or more maritime vessels, land vehicles and/or aircraft, each of which may be controlled by on-board personnel, or may be controlled by remote personnel and/or autonomously, such as an unmanned aerial vehicle (UAV) or unmanned surface vehicle (USV).

[0009] The mobile base stations are connected directly or indirectly (e.g. via one or more other mobile base stations) to one or more gateway nodes 3 which provide connectivity to other communication networks, such as the Internet. In a specific example of a 5G mesh network, the mobile base stations 1 may be mobile Integrated Access Backhaul (IAB) nodes and the one or more gateway nodes 3 may be Donor Distributed Units (DUs).

[0010] Each regulatory region is governed by its own regulatory body which licenses one or more mobile network operators (MNOs) operating in one or more licensed frequency bands. Each region may be defined by a polygon or other shape overlayed on a map, with associated frequency bands and maximum equivalent isotropic radiated power (EIRP) configured to meet power flux density (PFD) limits within the regulatory region.

[0011] As shown in Fig. 1, a regulatory region may be defined by a regulatory border 4 beyond which a mobile base station 1 is not permitted to cause RF emissions 5 exceeding defined limits,

such as PFD limits. In this figure, the emissions 5 are represented by one or more beam patterns emitted by the mobile base station 1. It is therefore desirable, and may be required for regulatory reasons, to control the RF emissions of the mobile base stations 1 so as to comply with such regulatory borders 4.

[0012] Fig. 2 shows a mesh network 10 interoperating with a satellite network 20 providing mobile satellite connectivity to the mobile base stations 1, in addition to the connectivity provided by the mesh network 10. The satellite network 20 and the mesh network 10 are both connected to a controller 6 which facilitates operation of the mesh network 10 using out-of-band connectivity via the satellite network 20, for example by interconnection between one or more ground stations 21 of the satellite network 20 and one or more gateways 3 of the mesh network 10, for example via a network 7 such as the Internet. In one example, the controller may be an Inmarsat software-defined network (SDN) controller (SDN-C) having connectivity to one or more Inmarsat satellite networks 20, such as the Broadband Global Area Network (BGAN) (TM) or Global Xpress (GX) (TM), in addition to one or more mesh networks 10 such as 5G and/or LTE networks. In at least some embodiments of the invention, the satellite network(s) 20 operate as underlay network(s) to provide additional functionality to assist regulatory compliance in the mesh network 10.

[0013] Regulatory compliance may be enforced on a mobile base station 1 located on a mobile platform 2 by defining one or more blockage zones 30, within the antenna coverage of the mobile base station 1, within which emissions are inhibited. The antenna 35 is aligned with the mobile platform so that the blockage zone(s) may be defined with reference to the orientation of the mobile platform. In the case of a maritime vessel, a datum point of the antenna may be aligned with the bow of the vessel and the blockage zone(s) may be defined with reference to the heading of the vessel. In the example shown in Fig. 3, a blockage zone is defined from approximately 60 degrees through to 180 degrees in azimuth relative to the direction in which the bow is pointing. This defines a blockage arc of approximately 120 degrees in which antenna emissions are inhibited. The blockage zone may additionally or alternatively be defined in terms of angle of elevation.

[0014] As shown in Figure 4, the antenna 35 may generate an adjustable beam pattern 36 of one or more directional beams 36, such as an array of beams covering different directions and/or one or more steerable beams. The beams may be steered and/or selectively inhibited according to a beam mask.

[0015] Some or all of the following functionalities may be implemented in the controller 6, or one or more nodes connected to the controller 6, in order to define the blockage zones of the mobile base stations 1:

- A positional data feed, preferably in real time or near real time, of the locations and bearings of the mobile platforms 2. The positional data may be obtained from the mobile platforms 2 e.g. from GNSS position and digital compass readings obtained on board the mobile platforms 2.
- A data model of each region within which the mesh network 10 is enabled, including for example a regional map overlayed with the positions of the mobile platforms 2, the gateway node 3 positions and the regulatory borders 4. Additional features, such as obstructions that may affect RF emissions, may also be included.
- A model of the network topology, which may be updated in real time or near real-time
- A function to calculate a new beam mask and corresponding antenna blockage zone 30 for a mobile platform 2 each time its position (e.g. location and bearing) is updated.

[0016] The controller 6 may provide a control plane available to all of the available networks so as to provide in- and out-of-band control. The controller may use an out-of-band network (e.g. satellite network(s) 20) to create a management plane that offers dynamic configuration of a mobile base station 1. This facilitates when and how the mobile base station 1 joins the network and provides a means to make topology adaptations. It also caters for adjustments to parameters that control beam masking, frequency allocations, maximum EIRP etc.

[0017] The out of band network is implemented as a management network that is made available over multiple bearers. This may be implemented using a software-defined wide area network (SD-WAN), as shown for example in Figure 5. The mobile platform 2 (e.g. maritime vessel) includes an SD-WAN end point 41 having respective data plane connections to the mesh network 10 and to one or more out of band networks, for example via:

- an IAB unit 42 for connection to the mesh network, including the gateway;
- a GX terminal 43 for connection to a GX network i.e. via a GX satellite 22 to a GX ground station 23; and
- a BGAN terminal 44 for connection to a BGAN network i.e. via a BGAN satellite 22 to a BGAN ground station 23.

[0018] The SD-WAN end point is connected by a management network within the mobile platform, for example via a LAN bridge 45 to an O1 termination 46 which is connected via a local management link to the IAB unit 42.

[0019] The mesh network 10 and each of the out of band networks 20 are connected to a common SD-WAN terminator node 50 which is connected via a core network 51 and a management network virtual routing and forwarding (VRF) function 52 to the controller 6, in this case the SDN-C.

[0020] The above provides an end-to-end management network that is made available seamlessly across multiple mobile communications networks 10, 20 (e.g. 5G mesh, GX and BGAN) and which may implement an SD-WAN policy to route traffic across the best available communication network. This ensures that the SDN-C management plane has the resilience necessary for reliable operation.

[0021] The SDN-C 6 maintains a data model that represents the current state of the network, including where each IAB node is located, its orientation and any parent nodes. The SDN-C 6 detects when a mobile platform 2 enters a region, for example via positional data updates encapsulated within VES notifications and calculates a list of possible parent nodes for the associated IAB node 1. A parent node can be a gateway 3 (in the 5G case a Donor DU) or another mobile node 1 that is already in the network. A parent node becomes essentially the next upstream hop in the mesh network 10 with ability to backhaul user traffic towards a Donor DU acting as a base station and gateway 3.

[0022] Figure 6 is a flowchart of a method of determining dynamic beam masking in an embodiment of the invention, as performed for example by the controller 6 (e.g. SDN-C) and/or nodes connected thereto. If a new current position for a mobile platform 2 is detected (S1), at least two parallel processes are initiated. In the first process, the overall network topology is updated (S2) and stored (S3). Next, it is determined whether the mobile node associated with the mobile platform has any child nodes within the mesh network (S4). If so, the DU is enabled for that mobile node (S5); otherwise, it is disabled (S6). In the second process, a beam mask is determined (S7) based on the regional and regulatory map (S8).

[0023] The first and second processes terminate in a process to update the configuration of the beam mask (S9) for the mobile node 1. If the configuration has changed (S10), then the new configuration is applied to the mobile node e.g. by changing the blockage zone(s) applied to the antenna(s) of the mobile node. Otherwise, no change is made (S11). In either case, the method then restarts.

[0024] Depending on the cadence of positional updates and the latency incurred when processing data, applying a new beam mask could take quite some time. This means there will be a level of uncertainty with respect to the actual vessel position and whether a beam mask adequately protects the regulatory zones.

[0025] To address the above problem, a predicted mobile platform position at the next position interval can be derived and an associated beam mask calculated. These two beam masks can then be summed together to create a wider and more conservative blockage arc. This technique ensures a higher level of confidence that the mobile node will not emit past a regulatory border. Alternatively or additionally, a plurality of beam masks may be calculated, each for a corresponding different time, and transmitted to the mobile node for use at the corresponding time if no update is received from the controller 6.

[0026] The predicted mobile platform position may be derived from proposed and/or historical route data e.g. from a flight plan for an aircraft or a route plan for a ship. The route data plans may be provided by operators of the mobile platforms or may be deduced e.g. from standard shipping or aircraft routes.

[0027] Figure 7 shows a variant of the method of Figure 6, including a third parallel process to make use of predicted mobile platform position. In this third process, when a new current position is detected for the mobile platform, the positional data is stored in a database (S12). Next, previous positions for the mobile platform are read from the database (S13) e.g. the last N readings where N is a predetermined number, or the readings over a previous predetermined time period, to produce historical position information (S14). This historical position information is then used to predict the next position of the mobile platform (S15. S16), which is then used to calculate the next beam mask (S17), based on the stored regional and regulatory map (S8). The current and next beam masks are then summed together (S18) as described above.

[0028] Figure 8 shows an example use case of calculation of a beam mask for a mobile node 1 within emission range of a regulatory border. The permitted emission direction range 60 is shown as a shaded azimuthal range. As the mobile platform 2 moves between different positions, the angle of the permitted emission range 60 relative to the bearing of the mobile platform 2 changes so as to avoid emissions towards the regulatory border 4.

[0029] The SDN-C 6 may access information about the radio performance of the IAB units 42 and the agreed PFD limits both within and outside the regulatory region, and may determine a maximum range using a maximum EIRP in ideal conditions at which the radio emissions fall below the agreed limits for outside of the regulatory zone. Using the maximum range distance, a circle 61 may be defined around an IAB node 1 to indicate the point at which emissions fall within PFD limits. The points at which the circle 61 intersects with regulatory borders 4 may be used to define a beam masking arc 30.

[0030] Figure 9 shows one example of calculation of a beam masking arc 30, in which a mobile platform 2 is within emission range of a regulatory border 4 represented by a single line. The

shaded circle 61 shows the range within which radio emissions from the mobile IAB unit 42 are determined as above the PFD limits. A line may be drawn from the centre of the mobile IAB unit 42 to each of the two points where the circle 61 intersects a point on the regulatory border 4. The angle between the two lines defines a blockage arc 30 within which the beam area must be masked; in this example, the blockage arc is formed from 95° to 185° relative to the heading of the mobile platform 2.

[0031] Figure 10 shows another example in which the mobile platform 2 is within emission range of a regulatory border 4 represented by two lines at an angle. In this case, the circle intersects the regulatory border at four points and two separate beam masks are derived, one extending from 95° to 185° and another from 205° to 260° relative to the heading of the mobile platform 2. Emissions between 185° and 205° are not masked.

[0032] Figure 11 shows an example of a computer system 1000 on which one or more functionalities of embodiments of the invention may be implemented as computer-readable code. For example, functions of the controller 6 and/or mobile nodes 1 described above can be implemented in system 1000. After reading this description, it will become apparent to a person skilled in the relevant art how to implement the invention using other computer systems and/or computer architectures.

[0033] Computer system 1000 includes one or more processors, such as processor 1004. Processor 1004 can be a special purpose or a general-purpose processor. Processor 1004 is connected to a communication infrastructure 1006 (for example, a bus or network). Computer system 1000 may include a user input interface 1003 connected to one or more input device(s) 1005 and an output interface 1007 connected to one or more output devices 1009, which may include a display. Input devices 1005 may include, for example, a connected device such as a mouse or touchpad, a keyboard, a touchscreen such as a resistive or capacitive touchscreen, etc.

[0034] Computer system 1000 also includes a main memory 1008, preferably random-access memory (RAM), and may also include a secondary memory 1010. Secondary memory 1010 may include, for example, a hard disk drive, a removable storage drive, flash memory, a memory stick, and/or any similar non-volatile storage mechanism. As will be appreciated by persons skilled in the relevant art(s), removable storage may include a non-transitory computer usable storage medium having stored therein computer software and/or data.

[0035] In alternative implementations, secondary memory 1010 may include other similar means for allowing computer programs or other instructions to be loaded into computer system 1000. Such means may include, for example, a removable storage unit and an interface which allow software and data to be transferred from the removable storage unit to computer system 1000.

[0036] Computer system 1000 may also include a communications interface 1024 implemented for example at the operating system level to allow data to be transferred between computer system 1000 and external devices, for example as signals over a communication channel. Communications interface 1024 may include a modem, a network interface (such as an Ethernet card), a communications port, a PCMCIA slot and card, or the like.

[0037] Various embodiments of the present invention may be implemented by software and/or firmware (also called computer programs, instructions or computer control logic) to program programmable hardware, or hardware including special-purpose hardwired circuits such as application-specific integrated circuits (ASICs), programmable logic devices (PLDs), field-programmable gate arrays (FPGAs), etc. of the computer system 1000, or a combination thereof.

[0038] Computer programs for use in implementing the techniques introduced here may be stored on a machine-readable storage medium and may be executed by one or more general-purpose or special-purpose programmable microprocessors. The terms “computer program medium”, “non-transitory computer readable medium” and “computer usable medium” introduced herein can generally refer to media such as removable storage unit, removable storage unit, and a hard disk installed in hard disk drive. Computer program medium, computer readable storage medium, and computer usable medium can also refer to memories, such as main memory 1008 and secondary memory 1010, which can be memory semiconductors (e.g. DRAMs, etc.). These computer program products are means for providing software to computer system 1000.

[0103] Computer programs are stored in main memory 1008 and/or secondary memory 1010. Computer programs may also be received via communications interface 1024. Such computer programs, when executed, enable computer system 1000 to implement the present invention as described herein. In particular, the computer programs, when executed, enable processor 1004 to implement the processes of embodiments of the present invention as described above. Accordingly, such computer programs represent controllers of the computer system 1000. Where the invention is implemented using software, the software may be stored in a computer program product and loaded into computer system 1000 using removable storage or communications interface 1024.

Alternative Embodiments

[0039] Alternative embodiments, which may become apparent on reading the above description, may nevertheless fall within the scope of the invention as defined by the accompanying claims.

[0040] Although the specific embodiment above addresses regulatory compliance with respect to emissions across regulatory borders, other embodiments may implement other types of geographical restriction requirements on emissions.

Claims

1. A method of controlling emissions from a mobile node in a mesh network, the method comprising, at a controller remote from the mobile node:

obtaining positional data relating to the mobile node;

comparing the positional data with a geographical emissions restriction requirement;

determining a directional beam mask for the mobile node so as to meet the geographical emissions restriction requirement; and

controlling the emissions of the mobile node based on the beam mask;

wherein the controller has out of band connectivity with the mobile node via one or more alternative networks other than the mesh network.

2. The method of claim 1, wherein the controller controls the mobile node via the one or more alternative networks so as to apply the beam mask to the emissions of the mobile node.

3. The method of any preceding claim, wherein the controller controls connectivity of the mobile node to the mesh network via the one or more alternative networks.

4. The method of any preceding claim, wherein said positional data is obtained at least in part via said one of more alternative networks.

5. The method of any preceding claim, wherein the beam mask is determined by determining a maximum range at which emissions from the mobile node may exceed a limit of a regulatory emissions boundary.

6. The method of claim 5, including determining at least two points at which the regulatory emissions boundary is at said maximum range from the mobile node, and determining an angle between said at least two points within which the beam mask is applied.

7. The method of any preceding claim, including deriving a future position of the mobile node, and determining the beam mask based on said positional data and said future position.

8. The method of claim 7, wherein the beam mask is determined by summing a beam mask based on the positional data and a beam mask based on the derived future position.

9. The method of claim 7, wherein a plurality of beam masks are determined for different corresponding future times and are transmitted to the mobile node for use at said corresponding future times.

10. The method of any preceding claim, wherein the positional data includes data indicating orientation of a mobile platform of the mobile node, and the beam mask is defined relative to said orientation.

11. A controller arranged to perform the method of any preceding claim.

12. A computer-readable storage medium comprising instructions which, when executed by a computer, cause the computer to carry out the method of any one of claims 1 to 10.