(12) UK Patent Application (19) GB (11) 2512748 (13) A

(43) Date of A Publication 08.10.2014

(21) Application No: 1405791.3

(22) Date of Filing: 31.03.2014

(30) Priority Data:
(31) 1403312 (32) 25.02.2014 (33) GB
(31) 1403314 (32) 25.02.2014 (33) GB

(71) Applicant(s):
Cambridge Silicon Radio Limited
(Incorporated in the United Kingdom)
Churchill House, Cambridge Business Park,
Cowley Road, CAMBRIDGE, Cambridgeshire,
CB4 0WZ, United Kingdom

(72) Inventor(s):
Hugo Mark Tyson
Nicholas Guy Albert Graube

(74) Agent and/or Address for Service:
Slingsby Partners LLP
5 Chancery Lane, LONDON, EC4A 1BL,
United Kingdom

(51) INT CL:
H04W 52/02 (2009.01) ***H04W 84/18*** (2009.01)

(56) Documents Cited:
US 20110053493 A1 **US 20090232037 A1**
US 20090032037 A1

(58) Field of Search:
INT CL **H04W**
Other: **EPODOC, WPI**

(54) Title of the Invention: **Auto-configuration of a mesh relay's TX/RX schedule**
Abstract Title: **Power saving at a node in a wireless mesh network**

(57) A mesh network consists of wireless nodes which are each capable of relaying packets to their intended destination. Flood routing may be used to propagate the packets from source to destination via multiple hops. This implies that each device continuously listens for signals from other nodes within the network, which increases power consumption. To reduce power consumption, a wireless node detects the number of wireless nodes nearby which are also able to relay packets and adjusts its power consumption accordingly. For instance, if a large number of relay nodes is detected, the wireless node will conserve power by adjusting its listening schedule (e.g. increasing the time between each listening period), and vice versa. Alternatively, the wireless node may save power by reducing a receive range if a large number of nodes is detected. This may also reduce loading on a network due to less duplicate packets being relayed concurrently.

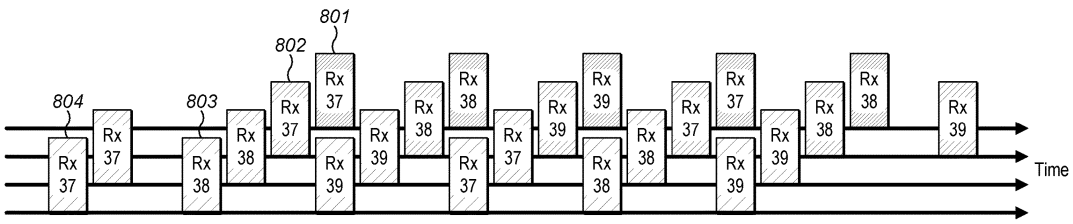

FIG. 8

GB 2512748 A

At least one drawing originally filed was informal and the print reproduced here is taken from a later filed formal copy.

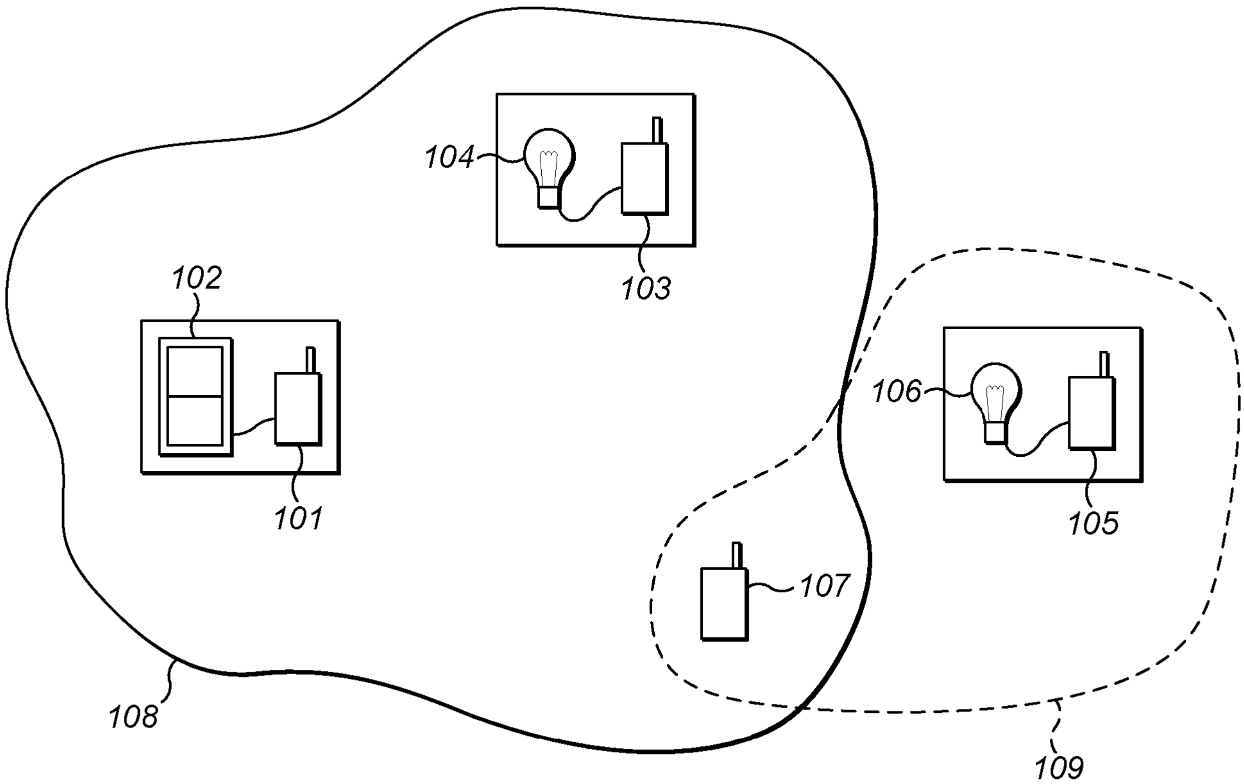

FIG. 1

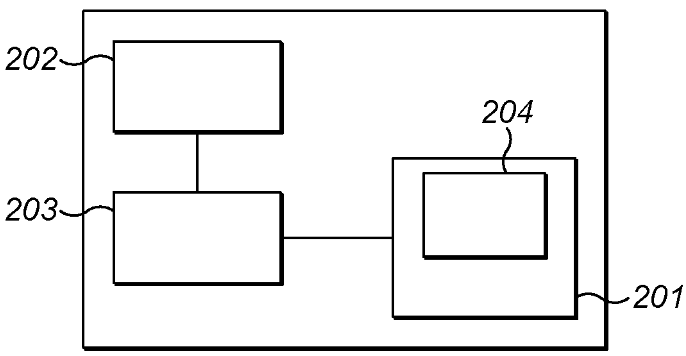

FIG. 2

22 04 14

22 04 14

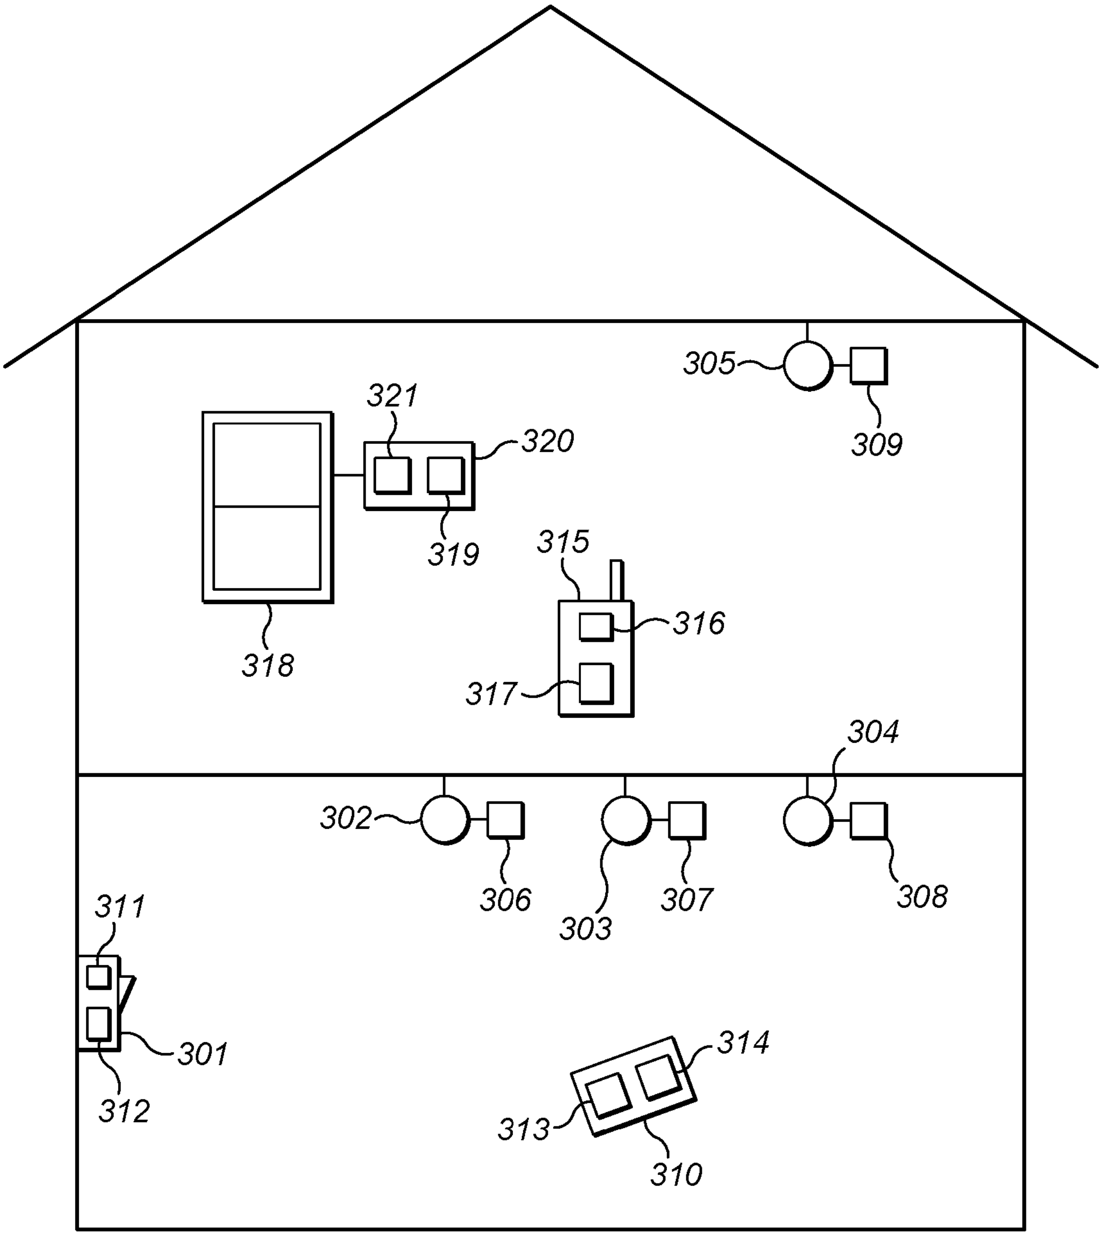
305
309
321
320
319
318
315
316
317
304
302
306
303
307
308
311
301
312
314
313
310

FIG. 3

22 04 14

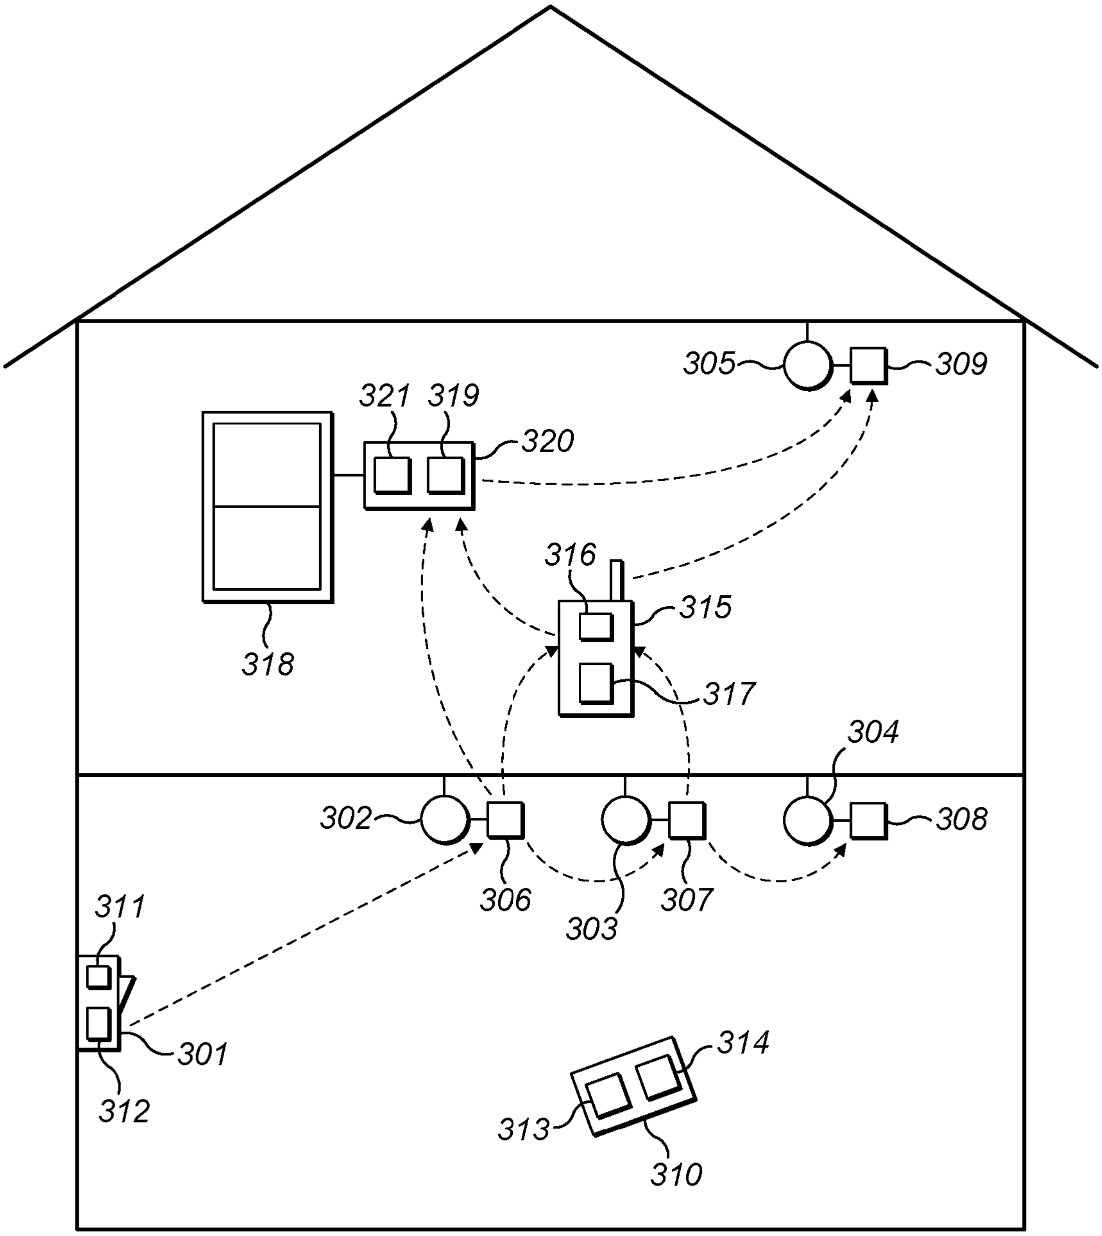

FIG. 4a

22 04 14

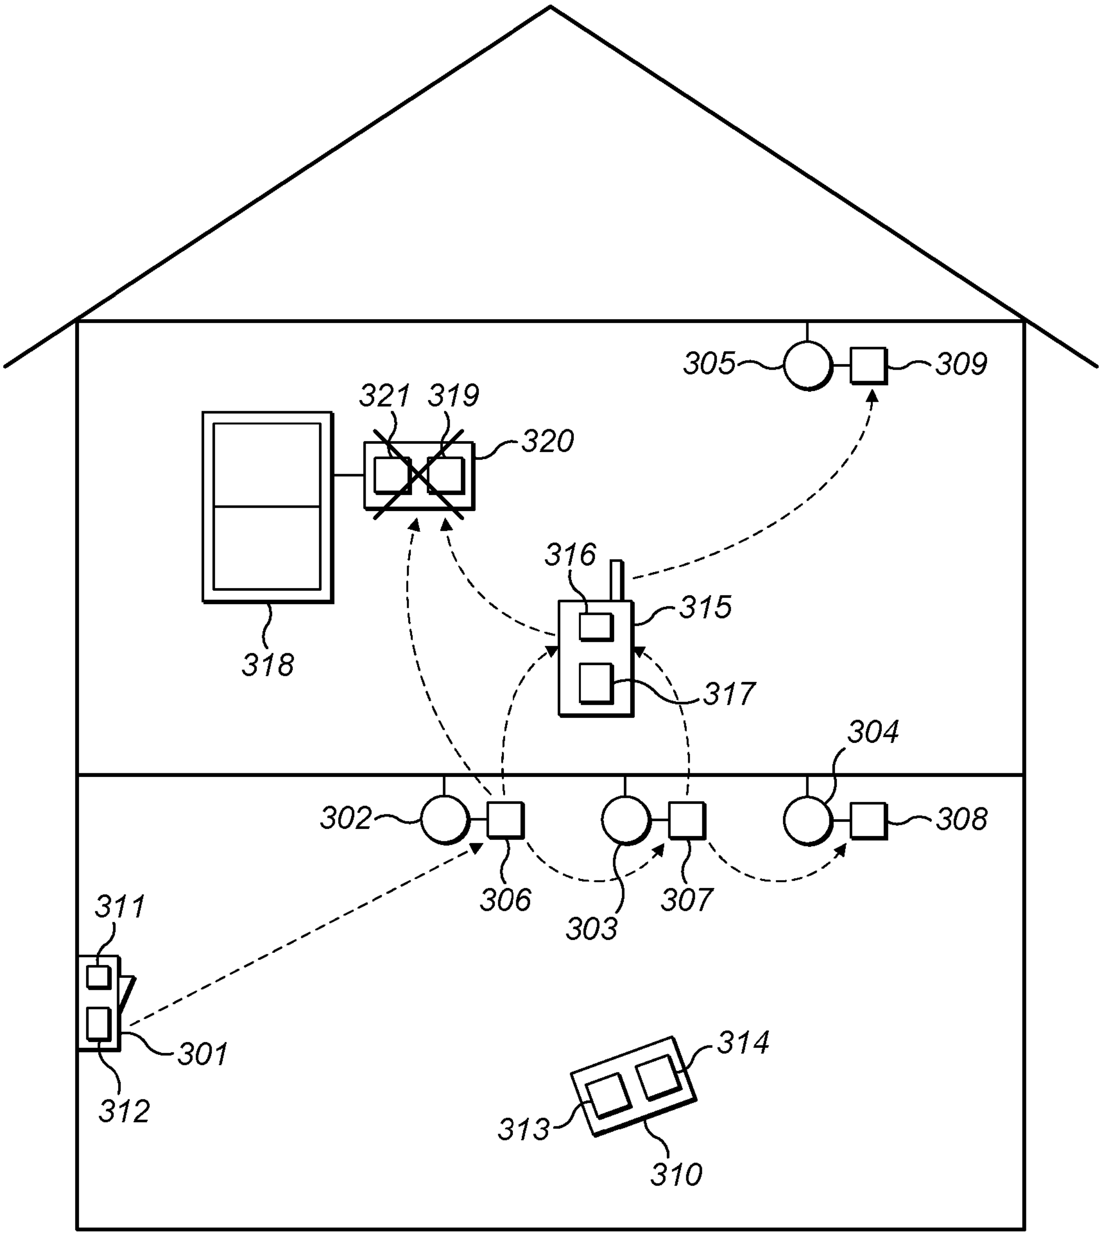

FIG. 4b

22 04 14

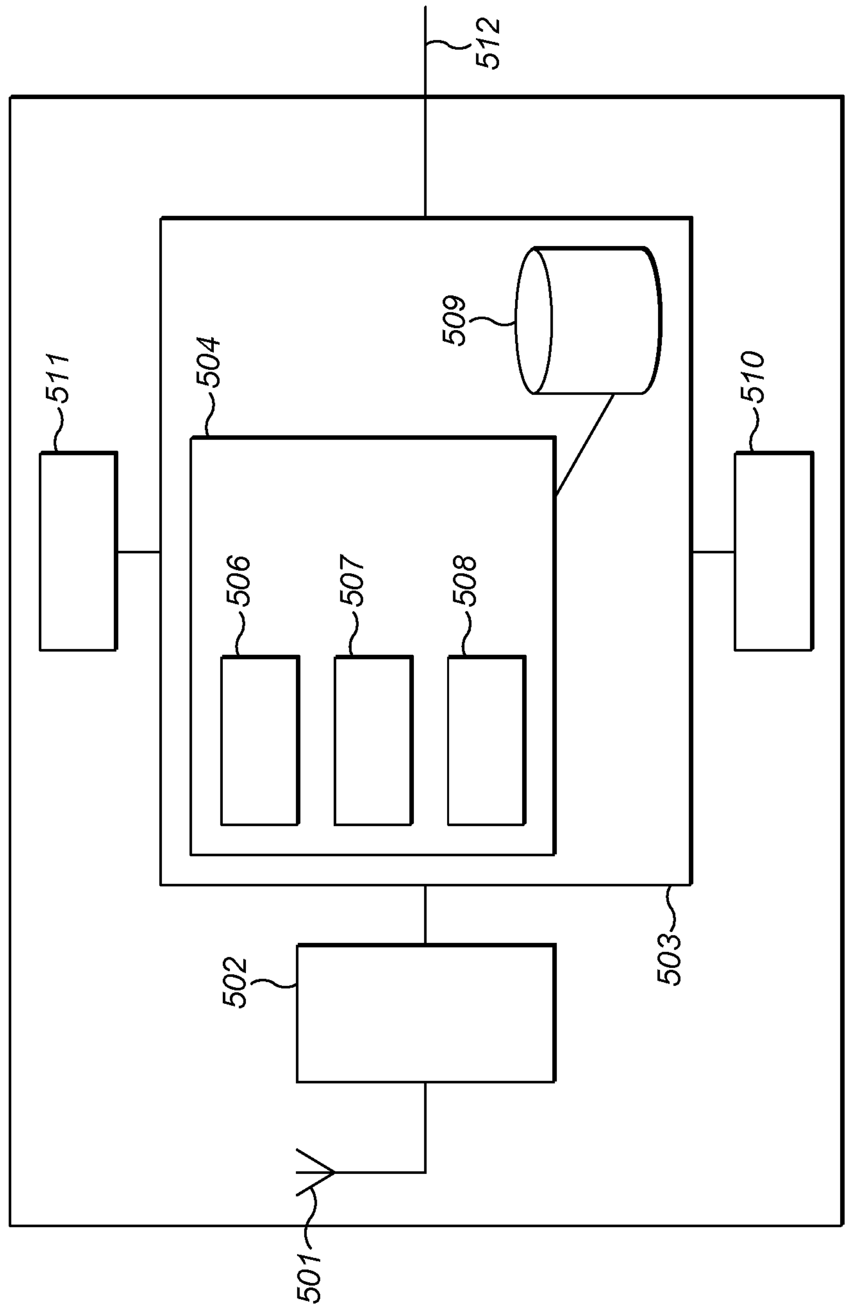

FIG. 5

22 04 14

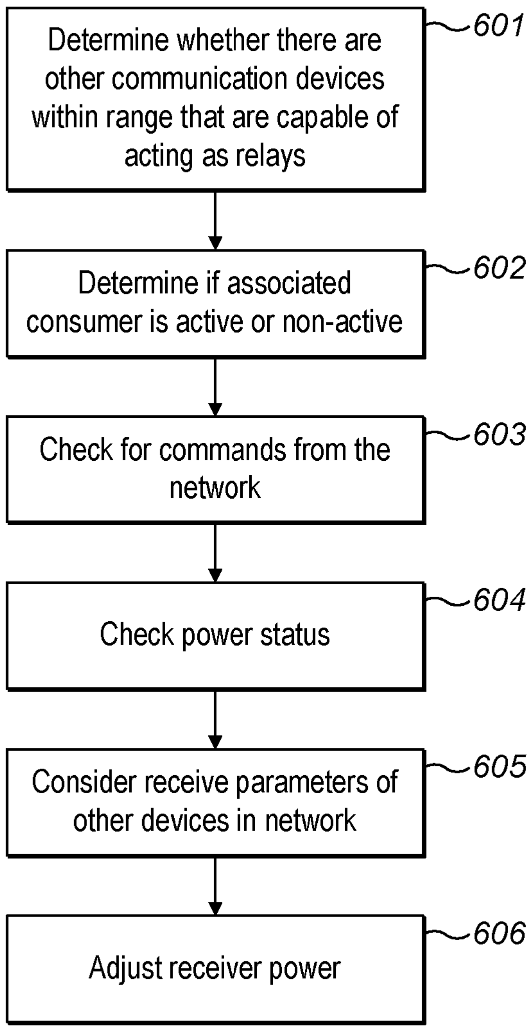
Determine whether there are other communication devices within range that are capable of acting as relays
601
Determine if associated consumer is active or non-active
602
Check for commands from the network
603
Check power status
604
Consider receive parameters of other devices in network
605
Adjust receiver power
606

FIG. 6

22 04 14

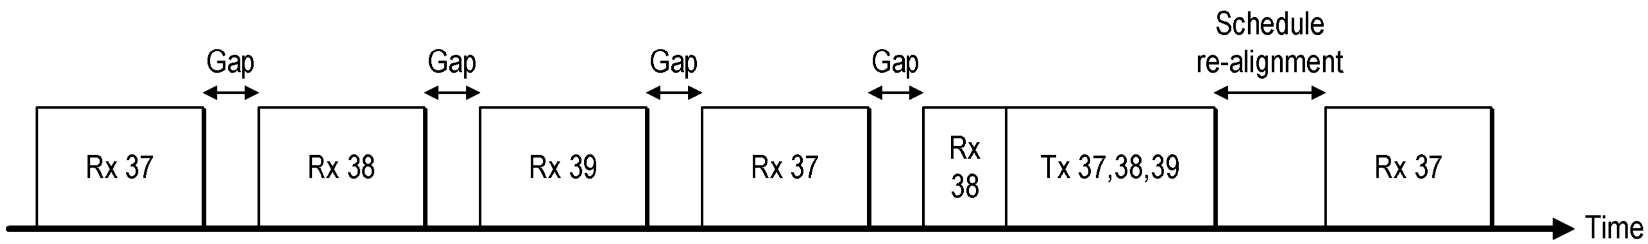

FIG. 7a

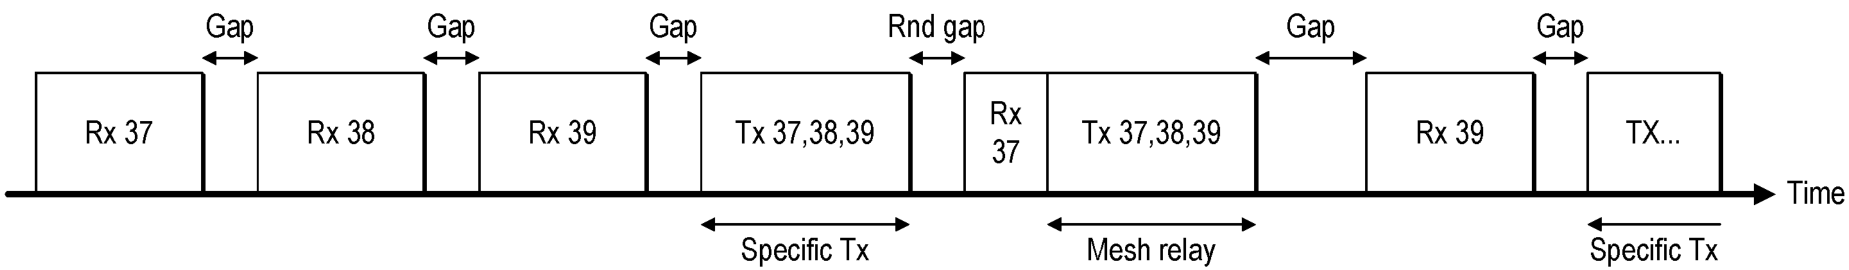

FIG. 7b

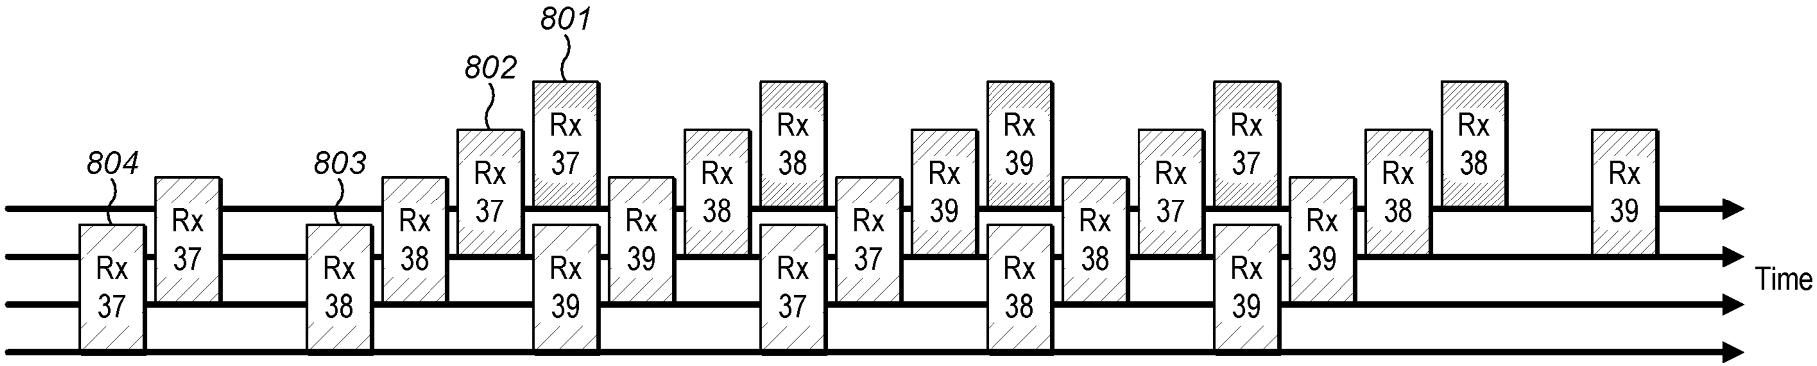

FIG. 8

AUTO-CONFIGURATION OF A MESH RELAY'S TX/RX SCHEDULE

This invention relates to a communication device that is capable of relaying packets in a mesh network.

Figure 1 shows a distributed network. The network comprises a number of communication devices 101, 103, 105 and 107. Each device can communicate wirelessly with the other devices that are in effective range of it. In this example the network is a mesh network, meaning that each device relays data for the network. The devices can communicate to propagate data between them. For example, if device 101 transmits a signal, that signal can be received by devices 103 and 107 which are within range of device 101. Devices 103 and 107 can then relay the signal received from device 101 so that it can be received by device 105, which is out of range of device 101. The coverage area of device 101 is illustrated at 108 and the coverage area of device 105 is illustrated at 109. This method of communication allows devices to communicate even though they are out of direct range or not synchronised with each other. Each device may also be connected to, or integrated within, an associated consumer device. So device 101 is connected to a sensor that detects whether window 102 is open or closed, and devices 103 and 105 are connected to light fittings 104 and 106 respectively.

Many mesh networks send data using complex routing tables. The routing tables store routes from one network device to another so that messages can be propagated from source to destination via a series of hops. The topology of the network has to be known in order that routes between the various devices can be determined and stored. An alternative is flood routing. In this method messages do not travel from one device to another via a predefined route. Instead messages are broadcast and any device in range that receives a message retransmits it. A message thus propagates its way through the network, potentially reaching its destination via a number of different routes. Flood routing is very simple to implement and although it may appear inefficient has a number of advantages, particularly for ad hoc networks that may change their topology on a random basis.

Flood routing implies that all devices should listen continuously for signals from other devices in the network, otherwise there is a risk that data might not reach its destination. Continuous listening increases power consumption. In current mesh networks this is often unimportant because most mesh devices have access to a mains power source, which eliminates the requirement for power saving. It does, however, limit the range of devices that can form part of the network. There is a need to open up mesh networks to a wider range of devices, including those with severe power restrictions.

According to one embodiment, there is provided a communication device capable of communicating over a communication network that is configured such transport of packets through the network is provided by each communication device in the network listening for and relaying packets, the communication device comprising a relay unit configured to listen for packets and relay them over the network, a detection unit configured to identify other communication devices in the network that are also capable of acting as relay devices and a power controller configured to adjust the power that the communication device uses to listen for packets in dependence on the other devices identified by the detection unit.

The power controller may be configured adjust the power in dependence on the number of other devices identified by the detection unit.

The power controller may be configured adjust the power in dependence on receive parameters associated with the other devices identified by the detection unit.

The receive parameters may include one or more of a listening schedule, a power status and a receive range.

The power controller may be configured adjust the power in dependence on a mode of operation of a consumer device associated with the communication device.

The power controller may be configured adjust the power in dependence on whether its associated consumer device is active or non-active with respect to the network.

The power controller may be configured adjust the power in dependence on information received from the network about the current transport capabilities of the network.

The power controller may be configured to adjust the power in dependence on a power status associated with the communication device.

The power controller may be configured to control the rate at which it adjusts the power in dependence on a power status associated with the communication device.

The detection unit may be configured to identify the other devices using packets that have previously been relayed to it over the network.

The detection unit may be configured to identify the other devices using packets that it has been relayed multiple times.

The detection unit may be configured to identify the other devices using a lifetime value incorporated in packets that have previously been relayed to it over the network.

The power controller may be configured to adjust the power by controlling the amount of time that the relay unit listens for packets.

The power controller may be configured to control a period of time for which the relay unit listens continuously for packets.

The power controller may be configured to control a time interval between periods of time for which the relay unit listens continuously for packets.

The power controller may be configured to control the relay unit such that, the higher the number of devices identified by the detection unit, the less time that the relay unit listens for packets.

The power controller may be configured to adjust the power by controlling a receive range of the communication device.

The power controller may be configured to adjust the power by controlling the usage of receive circuitry in the communication device.

The relay unit may be configured to listen to packets asynchronously from other devices in the network.

According to a second embodiment, there is provided a method of communicating over a communication network that is configured such transport of packets through the network is provided by each communication device in the network listening for and relaying packets, the method comprising a communication device identifying other communication devices in the network that are also capable of acting as relay devices and adjusting the power that the communication device uses to listen for packets in dependence on the other devices identified by the detection unit.

The present invention will now be described by way of example with reference to the accompanying drawings. In the drawings:

Figure 1 illustrates a distributed network;

Figure 2 illustrates a wireless communications device;

Figure 3 illustrates a distributed lighting system installed in a house;

Figures 4a and 4b illustrate the transport of a packet through the distributed lighting system;

Figure 5 illustrates another example of a communication device;

Figure 6 illustrates an example of a method for controlling the listening windows of a communication device;

Figures 7a and 7b illustrate an example of a communication device listening for packets for the maximum time available; and

Figure 8 illustrates an example of the listening windows of four communication devices that have reduced the frequency of their listening.

In a wireless network data is transmitted via radio signals. In a wired implementation, data may be transmitted via electrical signals. Most commonly data will be arranged in “packets” incorporating payload data and some indication of the source device and the destination device. Packets take many different formats and the apparatus, network and method described herein are not limited to using a particular type of packet. The term “packet” is therefore used herein to denote any signal, data or message transmitted over the network.

There are three main reasons why it may be impractical for a communication device to spend significant amount of time listening for packets. The first is the cost in power. The second is the physical impossibility for many devices of receiving at the same time as transmitting. The third is the need to perform other functions outside of the network, which may not be possible while listening for packets. The challenge in a mesh network is to reconcile these drawbacks with each node’s responsibility to relay data for the network.

For many nodes in a mesh network it is simply impossible to listen for packets all the time. This is because the communication devices that implement those nodes are incapable of listening at the same time as they transmit. So, packets that arrive while the device is transmitting are inevitably lost. Thus it is more realistic to aim at maximising the number of packets received by a device than to aim at having every device receive all packets.

This leads to two further insights. First, it is often the case that a device acting as a relay in a mesh network does not actually need to receive all packets because its role is only to relay those packets, not to alter its behaviour in dependence on the data that those packets contain. Consequently, not only is it technically impossible for a single device to listen continuously, it is also (at least from the perspective of that particular device) often unnecessary. Second, from the perspective of the network it does not matter which particular device relays each packet; it only matters that the packet is relayed. Thus it may be possible to have one relay device reduce its receive power, for example by reducing its range or deliberately not listening for packets at all, on the basis that any gaps in its listening will be covered by other devices in the vicinity. It becomes possible to trade listening power against the number of relay devices. For example, one device listening for N seconds can be traded against two or more devices listening for K seconds, with $K < N$.

A general example of a communication device is shown in Figure 2. The communication device is capable of acting as a relay device in a mesh network. In most examples the communication device will be configured for wireless communication over the mesh network, although wired configurations are also possible. It comprises a relay unit 201, a detection unit 202 and a power controller 203. The relay unit may be configured to listen for packets. Typically those packets will be intended for one or more other devices in the network. The data content of the packets is thus immaterial to the communication device; its role is just to relay those packets by retransmitting them. The relay unit preferably either has access to a communication unit or incorporates a communication unit (not shown) so that it sees packets that have been received by the communication unit and can control those packets being retransmitted by the device. The detection unit may be configured to identify other devices in the network that are also acting as relay devices. There are a number of ways it might achieve this, and some examples are described below. The power controller may be configured to control the amount of power that the communication device uses to listen for packets to relay. Suitably the power controller controls the receive power in dependence on the other relay devices that are identified by the detection unit. The power controller may incorporate a timing unit 204 that is configured to adjust a time that the relay unit spends listening for packets.

Typically the higher the number of relay devices detected within radio range of the communication device, the less power it will use on average listening for packets. This may be achieved by reducing the length of the windows for which the communication device listens and/or by increasing the interval between windows. The communication device might alter the duty cycle of its receive operations in dependence on the number of other devices it detects within radio range. The communication device might reduce the power consumed by a listening operation, e.g. by reducing its reception range. Conversely if the communication device detects a decrease in the number of other relay devices within range in the network, it suitably increases the amount of power it spends listening to packets to compensate. The overall effect is that a communication device with a local mesh density that is relatively high, e.g. because it is in the middle of a network with many other devices within range, will tend to use less receive power than a communication device with a local mesh density that is relatively low, e.g. because it is on the edge of a network with few other devices within range.

Time that a communication device saves by not listening for packets may be spent performing other functions or in a low power mode to conserve battery life. Any gaps between the device’s listening windows are suitably filled by the other devices within radio range so that the performance of the network as a whole does not suffer. In fact, the performance of the network may improve as a consequence of fewer devices relaying the same packet. Communication devices within range of each other may coordinate their listening periods so at least one of them is listening at all times. An alternative is for the devices to be asynchronous so that largely continuous listening is achieved naturally due to random offsets between the device’s own listening periods.

One advantage of this approach is that it enables battery-powered relay devices become possible, whereas previously only mains-powered relay devices could realistically be considered. This is a significant step in permitting flexible mesh deployment and ad-hoc arrangements. It also enables any device configured to operate in accordance with the mesh protocol to form part of the mesh whilst still allowing other “primary” (i.e. non-mesh) functions to be performed.

The structures shown in Figure 2 (and indeed all block apparatus diagrams included herein) are intended to correspond to a number of functional blocks in an apparatus. This is for illustrative purposes only. Figure 2 is not intended to define a strict division between different parts of hardware on a chip or between different programs, procedures or functions in software. In some embodiments, some or all of the algorithms described herein may be performed wholly or partly in hardware. In many implementations, at least part of the relay unit, sense unit and timing unit may be implemented by a processor acting under software control (e.g. the CPU of a communication device). Any such software is preferably stored on a non-transient computer readable medium, such as a memory (RAM, cache, hard disk etc) or other storage means (USB stick, CD, disk etc).

An example of a network is shown in Figure 3, which represents a house having a distributed lighting system. The system comprises a light switch unit 301 and light fittings 302, 303, 304, 305. Light switch unit 301 is integrated with a wireless communication device 313. Light fittings 302 to 305 are integrated with respective wireless communication devices 306, 307, 308, 309. The house has a mains electrical supply which powers the light fittings and their respective wireless communication devices 306 to 309. Light switch unit 301 and its wireless communication device 313 are powered by a local battery 311.

The house contains other items of equipment that contain other wireless communication devices. For example, there is a tablet computer 310 which contains a wireless communication device 313, and a mobile phone 315 which contains a wireless communication device 316. There is also a sensor 330 for detecting the open/closed state of window 318, which contains communication device 319. Computer 310, phone 315 and sensor 330 are powered by batteries 314, 317 and 331 respectively.

Wireless communication devices 306 to 309, 310, 315 and 319 operate according to the same wireless communication protocol. That could be a relatively short-range protocol. For example the effective range of each device could be less than 35m. That characteristic can permit the devices to use less power for transmitting and/or receiving than would be expected in a longer range protocol. The protocol could be one that imposes no common time-base at or below the transport level, or below the application or presentation levels. In other words, the devices in the network operate asynchronously of each other. That characteristic can reduce the devices' power consumption by reducing their for need accurate clocks running continuously. In one example, the devices could operate according to the Bluetooth protocol, specifically the Bluetooth Low Energy protocol. The devices could use other protocols, for instance IEEE 803.11.

Devices 306 to 309 are configured cooperatively in order that the light fittings 302 to 305 know to respond to signals from the light switch 301. This may be done by the devices 306 to 309 storing a common identification code in their respective non-volatile memories. The identification code may be stored in the light switch when it is manufactured, and stored in the light fittings at the time they are installed in the house. They may be stored in the light fittings by means of another device such as mobile phone 315 communicating with the wireless device of the light switch to read its identification code, and then communicating with the wireless devices of the light fittings to cause them to store that same identification code. This code may be a network key, and it may be used to sign all packets sent over the network.

The network is preferably also configured to implement flood routing, which is well suited to ad hoc networks. The phone 315 and the tablet computer 310 are both portable devices that change location within the network as a user picks them up and moves them. They may also occasionally leave the network and then reappear some time later. For example, when a user takes them out of range of the network by taking them out of the house and later returns them to the house. The network's topology is thus subject to random alteration.

In order to avoid packets being bounced around the network indefinitely, each packet suitably includes a lifetime field that defines the lifetime of the packet within the network. A communication device that receives the packet suitably checks whether the lifetime field is equal to a threshold value before retransmitting the packet. If the lifetime value is equal to the threshold, the communication device does not retransmit the packet. Otherwise the communication device does retransmit the packet. In one example the lifetime field is a Time-To-Live (TTL) field. This is a value in the packet that is suitably decremented each time that the packet is retransmitted. In one example the TTL value is decremented by one at each retransmission, with each communication device that receives the packets retransmitting it until the TTL value is decremented to zero. In another example the lifetime field is a Max Hop Count (MHC) field. In this example each communication device stores a threshold MHC value, which is a positive, non-zero number. The MHC value in each packet may be incremented by one each time that the packet is retransmitted, with each communication device that receives the packets retransmitting it until the MHC value reaches the device's stored MHC threshold.

The communication devices in Figure 3 are all connected to or fully integrated with another device – a "consumer" – on behalf of which the communication device transmits and receives packets over the network. In many cases the primary function of the consumer may have nothing to do with the network. Consumer devices have varying levels of complexity. In one example a consumer device might be a tablet computer; in another it might just be a clock configured to count down to an expiry date of some perishable goods. It also possible for the communication device to be a consumer itself. An example of such a scenario might be when a communication device uses X10, which is a protocol designed to support the integration of electronic devices within the home.

A connection between the communication device and its associated consumer may be wired or wireless. The communication device may be contained within the same housing as the consumer. In many implementations the consumer device might be fully integrated with the consumer; they might even share circuitry. Often the communication device will be implemented by a chip within the consumer. An example of this is communication device 316 within phone 315. In other implementations the communication device and the consumer may be separate devices that are connected together. For example, the communication device might be a BLE tag connected to a PC.

For the purposes of this document, the communication device is considered to be the combination of hardware and/or software that implements the protocol governing the network, thereby implementing the packet transport that enables the consumer to communicate over the network.

Each communication device may be capable of acting as a relay in the network. An example of this is shown in Figure 4a, which shows the same distributed lighting system as Figure 3. The network is configured as a mesh network so, at least in theory, all devices that are part of the network have a responsibility to act as relays. A relay device suitably retransmits any packet that it recognises as having originated from the network (e.g. because it has been signed using a network key). The relay device might also take steps to prevent old packets from being continuously bounced around the network, e.g. by only forwarding the packet if it is new and/or by decrementing a "time-to-live" value in the packet before forwarding it on. Figure 3a shows an example of the network operating according to traditional mesh principles. Light switch 301 transmits a packet addressed to all of devices 306 to 309 instructing light fittings 303 to 305 to switch on. This packet is propagated by all devices that receive it, eventually reaching light fitting 305, which is out of range of light switch 301, the source of the packet.

Although the arrangement shown in Figure 3a is effective at propagating packets to all devices in the network, constant listening is an expensive operation and should preferably be avoided in contexts where power availability is an issue. Although both device 316 and device 319 are capable as acting as relays, they are both battery powered and would prefer to reduce power consumption where possible. Therefore, one or both of those devices may deliberately reduce their receive power.

In Figure 4b, device 319 has detected that devices 306, 316 and 309 are all within range and are all capable of acting as relays. Device 319 is battery powered and is running low on power. It has therefore shortened its listening periods so that it only listens for 10% of the time. The network is unaffected, however, as the packet from light switch 301 is still relayed by device 316. The alternative arrangement is also possible, with device 319 listening and relaying packets instead of device 316.

Another example of a communication device is shown in Figure 5. In this example the communication device is configured for wireless communication. The device of figure 5 comprises an antenna 501, a radio frequency front end 502 and a baseband processor 503. The baseband processor comprises a microprocessor 504 and a non-volatile memory 509. The non-volatile memory 509 stores in non-transitory form program code that is executable by the microprocessor to cause the baseband processor to implement the communication protocol of the network. In this example the non-volatile memory 509 stores in non-transitory form program code that is executable by the microprocessor to implement the relay unit 506, the detection unit 507, and the power control unit 508.

The device also comprises a clock 510, which can be turned on or off by the microprocessor 504 in order to save power, and an external wired connection 512 for exchanging information with the device’s associated consumer. This information may include the sensing external events (e.g. the operation of an associated user interface device such as a switch) or issuing control signals to associated appliances (e.g. light fittings). The device also comprises a power source 511, which may be a battery. The device may also be mains-powered.

The RF front end 502 and the baseband processor could be implemented on one or more integrated circuits.

A communication device may take a wide range of different factors into account when deciding whether (and how) to adjust its listening operations to conserve power. Examples of the steps this might involve are shown in Figure 6.

In step 601 the communication device may determine the number of communication devices within range that are capable of acting as relays. This may be termed the “local mesh density”. There are a number of ways the communication device might do this. Some examples are described below.

The communication device suitably stores a record of packets it has received previously. The communication device may be configured to derive information about its local mesh density based on this record. For example, if the communication device has received the same packet multiple times, this may indicate that there are multiple devices acting as relays within range. Similarly, if those multiple copies of the same packet have similar values in their lifetime fields, this may tend to indicate that there are multiple relay-capable devices close to the communication device. If the communication device receives the same packet with very different lifetime value fields, this may be less indicative of a local mesh density but may indicate a good level of relay capability in the network as a whole.

The communication device may send probe packets to see how many copies of that packet it receives back from the network. Preferably the communication take steps to prevent probe packets from wasting resources by bouncing around the network. Preferably the probe packets are retransmitted only once. This may be achieved by the communication device setting the lifetime value of the probe packet to an appropriate value, e.g. by setting the TTL value to one. This has the further advantage of providing the communication device with local information about device density, since any copy of the probe packet it receives back must necessarily have been relayed by a device within radio range. In one example the communication device may incorporate a flag or similar in the probe packet that causes any device that receives it not only to retransmit the packet but also to include its own receive parameters, such as receive capabilities, receive range, power status, listening schedule etc. in the retransmitted packet.

In another example the communication device may be configured to operate over the network in accordance with a protocol that includes protocol, transport and bearer layers. Probe packets may be injected directly into the bearer layer, thus bypassing the protocol and transport layer. The protocol may define a host stack and a host controller stack. Bluetooth is an example of such an arrangement. BLE in particular places a significant amount of intelligence in the host controller so that the host only needs to be woken up when it needs to perform some action (the host is assumed to consume more energy than the host controller). In fact, in some implementations a communication device may not have a separate host at all, with the host controller (which is mostly implemented in software) performing all tasks that are usually associated with the host. It may be the host controller that injects the probe packet into the bearer layer In one example probe packets may be normal broadcast packets, e.g. broadcast packets according to the protocol that underlies the network (such as BLE). These normal broadcast packets may be different from the mesh packets that are normally transmitted over the network (for a start, they may not be signed by the network key), but they will still be monitored by all devices in the network. This also renders it possible for a device outside the mesh network to inject probe packets into the network. For example, a probing node, whose only purpose is to periodically broadcast data, might be outside of the mesh network itself but still inject probe packets into the mesh network.

The communication device may also be configured to more directly obtain information about neighbouring devices that are acting as relays. For example, the device may send packets requesting that any device that receives it responds with details about its own receive capabilities. This may include information such as power status, listening schedule, receive and/or transmit range, active status etc. These packets may be broadcast with a lifetime value that prevents them from being retransmitted.

For example, the packets could be transmitted with a TTL value of zero.

The communication device may also be configured, together with other devices in the network, to implement a more general feedback program in which they exchange information about receive strategies/capabilities so that the devices can coordinate their listening operations. Request and probe packets may form part of such a feedback program.

Having determined the local mesh density, the communication device may be in a position to adjust its receive power accordingly. There are other factors that the communication device may also take into account, examples of which are represented by optional steps 602 to 605.

In step 602 the communication device determines whether it or its associated consumer are “active” or “non-active” with respect to the network. In some embodiments this step may be performed before the communication device detects other devices in the network, since some devices may be configured not to reduce their receive power at all unless their associated consumer is non-active.

In general a device may considered to be “active” if it is waiting to receive a packet from the network that will cause it to adapt its behaviour. This “waiting” does not require the device to be positively anticipating a packet; it just refers to a state in which the consumer is configured to listen for a packet that might cause it to perform some operation (no matter how insignificant that operation might be). The communication device may include a mode unit configured to determine whether its associated consumer is “active” or “non-active”. It may make this determination based on information received from the consumer, e.g. a code or setting received at switch-on identifying the type of device that the consumer is and/or a status update each time that the consumer changes from active to non-active and vice versa.

The term “mesh device” may be used to refer to a communication device together with its associated consumer. A mesh device may fall into one of two categories:

- Mesh Active Device (MAD): Devices whose functional purpose includes obeying mesh commands and thus is required to receive all MESH packets to perform this function.
- Mesh Passive Transport (MPT): Devices whose primary purpose in the network is to implement a mesh transport. They will listen for mesh packets and forward new information. They can additionally offer other services (and commonly will since the consumer at least usually has at least one functional purpose other than communicating over a mesh network). They are, however, not required to receive to all mesh packets to perform this function.

An MAD is typically constrained in its scheduling as it is important for the effective realisation of its function that it receives all commands addressed to it. This implies that an MAD has to spend significant amount of time listening for potential commands. It is generally not advisable to attempt to reduce this listening time. Practically, this is a moot point as the MAD will usually be associated with some actuator, which implies access to an inexhaustible or rechargeable power source and thus removes the requirement for power saving.

An MPT is essentially stateless nature with respect to the transmitted packets. It does not affect the behaviour of an individual MPT if it has not observed a particular packet. This makes it possible for an MPT to have shorter and/or less frequent listening “windows” if there are sufficient other MPTs in the system to compensate.

A mesh device might be an MPT if its consumer component falls into any one of the following categories:

- a device might be of a type that is not sent commands by the network;
- a device might be of a type that does not adapt its behaviour in response to commands from the network;
- a device might be configured to receive commands at predetermined intervals, so that the rest of the time it is not waiting to receive anything;

- a device might be of a type that would normally expect to change its behaviour in dependence on a signal from the network but at present it has not got an assigned role within the network, so is not waiting to receive anything;
- a device might be of a type whose primary function is entirely separate from the network; it neither sends nor receives commands, although it can usefully relay them.

A communication device preferably listens for the maximum time available if it falls into the MAD category. Otherwise it risks missing packets that it actually needs to receive. A communication device may, however, safely reduce its listening time if it falls into the MPT category.

In step 603 the communication device may check whether it has received any take information or commands from the network that it should take into account when determining its receive power. For example, it can be envisaged that in some situations there might be an insufficient density of devices for all non-active devices to substantially reduce their listening time. Therefore, if one device (e.g. a controller) determines that packets are taking too long to reach their destination or are being regularly dropped, it may send out an instruction for all devices to increase their listening time accordingly. The device that sent the instruction might have determined that the network transport is not performing well enough based on the number of acknowledgments that it is receiving to its own packets and/or the length of time that those acknowledgements are taking to reach it.

The communication device may consider its own power status when deciding how to control its receive power (step 604). This factor may be particularly relevant when it comes to deciding whether or not to accede to a request from the network. For example, whether it has access to mains power or not and, it is battery powered, the amount of power the battery has left. If the device is battery powered with little power remaining, it may decide to continue minimising its receive power as much as possible notwithstanding the instruction from the network.

The communication device might also take into account any information it has about neighbouring devices that are capable of acting as relays (step 605). In particular it may compare any information it may have about a neighbouring device's power supply with its own power situation. For example, if the communication device knows from a previous exchange of status information with its neighbouring devices that one or more of them is mains powered, while the device itself is battery powered, it may determine that the neighbouring mains powered device is far better placed to carry the burden of listening for packets and maintain its own power-saving approach accordingly. The communication device may also consider any information it has about the receive arrangements of neighbouring devices, and particularly their listening schedules.

Finally, in step 606, the communication device may control its receive power. There are a number of different options available to a communication device for controlling its receive power, including the following examples:

- Reducing the amount of time that the communication device listens for packets. This might involve reducing the length of the windows during which the communication device continuously listens for packets or their frequency (e.g. by reducing the communication device's listening duty cycle).
- Using fewer receive circuits such as amplifiers and correlators. This may reduce the receive range of the communication device.

If a communication device decides that it should adjust its receive power, it may make any adjustments slowly, particularly if it is reducing its receive power, to allow other devices in the network time to adjust their own receive power to compensate. The communication device may control the speed at which the communication device ramps its receive power up or down in dependence on a variety of different factors, including a power level of the communication device, one or more receive parameters of its neighbouring devices, the time of day (some devices may have a schedule in which they do not listen at certain times of day) etc.

The principles described herein enable a network to transfer packets via a stochastic transport mechanism. This technique is well suited to ad hoc networks. The network is able to achieve reliable transport between two arbitrary devices via intermediary devices, which are not expected to handle communications reliably, because there is a sufficient density of devices in the network to compensate for devices that have reduced their receive power. Relaxing the constraint on the reliability of an individual device permits better power management, opening up the network to a wider range of devices. The mechanism is self-administering, in that each device makes its own decisions about how to adjust its receive power, but is still able to achieve network-wide reliability.

An example of how devices can compensate for one another is illustrated in Figures 7a, 7b and Figure 8.

An example of a device listening as much as it can is illustrated in Figure 7a and 7b for a mesh network that communicates using the three advertising channels specified by the BLE protocol. In Figure 7a the MPT scans successively on BLE advertising channels 37, 38, and 39 until it receives some new information. Figures 7a and 7b show one transmission over the network. A packet may also be relayed by retransmitting it multiple times over each of the relevant channels. Gaps between transitions from channel to channel are significantly shorter than the receive durations. Gaps between channels or between transmission groups may consist of a random element added to a fixed (minimum) duration to increase non-synchronisation with other beacons. The listening process is repeated until data is detected. If the data is considered to be a mesh packet that is valid and new (i.e. so that it is a packet for relaying), the receive operation is stopped for an immediate forwarding of the packet on all three channels. If the mesh packet is not for relaying, reception continues for its expected duration. Once the data has been transmitted, the standard schedule realigned to what it would have been if it had not been interrupted. Resuming the same schedule may be important to avoid devices accidentally synchronising after a relay operation. In figure 6b, a similar operation is followed but the device additionally performs a non-mesh transmit for additional function delivery. This non-mesh transmit and the relay transmit represent times when it is impossible for most current devices to listen. Thus, while the device tries to listen continuously, in practice this is not realisable so the device simply listens for the maximum time possible.

An alternative implementation is shown in Figure 8. Figure 8 shows the listening windows of four different MPTs. In this implementation devices 801, 802, 803 and 804 have all reduced their listening windows. The overall effect on transport in the network is negligible, however, since the times when one device is not listening are filled by another device. Thus it is possible to trade receiving time (and packet drop by individual devices) against an increase in the number of listening devices. Receive windows for individual devices are reduced, and compensated by the density of devices in the mesh.

The applicant hereby discloses in isolation each individual feature described herein and any combination of two or more such features, to the extent that such features or combinations are capable of being carried out based on the present specification as a whole in the light of the common general knowledge of a person skilled in the art, irrespective of whether such features or combinations of features solve any problems disclosed herein, and without limitation to the scope of the claims. The applicant indicates that aspects of the present invention may consist of any such individual feature or combination of features. In view of the foregoing description it will be evident to a person skilled in the art that various modifications may be made within the scope of the invention.

CLAIMS

1. A communication device capable of communicating over a communication network that is configured such transport of packets through the network is provided by each communication device in the network listening for and relaying packets, the communication device comprising:

a relay unit configured to listen for packets and relay them over the network;

a detection unit configured to identify other communication devices in the network that are also capable of acting as relay devices; and

a power controller configured to adjust the power that the communication device uses to listen for packets in dependence on the other devices identified by the detection unit.

2. A communication network as claimed in claim 1, the power controller being configured adjust the power in dependence on the number of other devices identified by the detection unit.

3. A communication network as claimed in any preceding claim, the power controller being configured adjust the power in dependence on receive parameters associated with the other devices identified by the detection unit.

4. A communication network as claimed in claim 3, in which the receive parameters include one or more of a listening schedule, a power status and a receive range.

5. A communication network as claimed in any preceding claim, the power controller being configured adjust the power in dependence on a mode of operation of a consumer device associated with the communication device.

6. A communication network as claimed in claim 5, the power controller being configured adjust the power in dependence on whether its associated consumer device is active or non-active with respect to the network.

7. A communication network as claimed in any preceding claim, the power controller being configured adjust the power in dependence on information received from the network about the current transport capabilities of the network.

8. A communication network as claimed in any preceding claim, the power controller being configured to adjust the power in dependence on a power status associated with the communication device.

9. A communication network as claimed in any preceding claim, the power controller being configured to control the rate at which it adjusts the power in dependence on a power status associated with the communication device.

10. A communication network as claimed in any preceding claim, the detection unit being configured to identify the other devices using packets that have previously been relayed to it over the network.

11. A communication network as claimed in claim 10, the detection unit being configured to identify the other devices using packets that it has been relayed multiple times.

12. A communication network as claimed in claim 10 or 11, the detection unit being configured to identify the other devices using a lifetime value incorporated in packets that have previously been relayed to it over the network.

13. A communication network as claimed in any preceding claim, the power controller being configured to adjust the power by controlling the amount of time that the relay unit listens for packets.

14. A communication network as claimed in claim 13, the power controller being configured to control a period of time for which the relay unit listens continuously for packets.

15. A communication network as claimed in claim 13 or 14, the power controller being configured to control a time interval between periods of time for which the relay unit listens continuously for packets.

16. A communication network as claimed in any preceding claim, the power controller being configured to control the relay unit such that, the higher the number of devices identified by the detection unit, the less time that the relay unit listens for packets.

17. A communication network as claimed in any preceding claim, the power controller being configured to adjust the power by controlling a receive range of the communication device.

18. A communication network as claimed in any preceding claim, the power controller being configured to adjust the power by controlling the usage of receive circuitry in the communication device.

19. A communication network as claimed in any preceding claim, the relay unit being configured to listen to packets asynchronously from other devices in the network.

20. A method of communicating over a communication network that is configured such transport of packets through the network is provided by each communication device in the network listening for and relaying packets, the method comprising a communication device:

identifying other communication devices in the network that are also capable of acting as relay devices; and

adjusting the power that the communication device uses to listen for packets in dependence on the other devices identified by the detection unit.

21. A communication device substantially as herein described with reference to the accompanying claims.

22. A communication network substantially as herein described with reference to the accompanying claims.

23. A method substantially as herein described with reference to the accompanying claims.

Intellectual Property Office

Application No:	GB1405791.3	**Examiner:**	Mr Steve Evans
Claims searched:	All	**Date of search:**	25 June 2014

Patents Act 1977: Search Report under Section 17

Documents considered to be relevant:

Category	Relevant to claims	Identity of document and passage or figure of particular relevance
X	1-5 & 20	US 2009/32037 A1 (NOKIA SIEMEND NETWORKS) - Whole document see especially paragraphs 31 and 34
X	1-5 & 20	US 2011/053493 A1 (OKI ELECTRIC) - Whole document

Categories:

X	Document indicating lack of novelty or inventive step	A	Document indicating technological background and/or state of the art.
Y	Document indicating lack of inventive step if combined with one or more other documents of same category.	P	Document published on or after the declared priority date but before the filing date of this invention.
&	Member of the same patent family	E	Patent document published on or after, but with priority date earlier than, the filing date of this application.

Field of Search:

Search of GB, EP, WO & US patent documents classified in the following areas of the UKC^{X} :

Worldwide search of patent documents classified in the following areas of the IPC

H04W

The following online and other databases have been used in the preparation of this search report

EPODOC, WPI

International Classification:

Subclass	Subgroup	Valid From
None		

Intellectual Property Office

Application No: GB1405791.3 **Examiner:** Mr Steve Evans

Claims searched: All **Date of search:** 25 June 2014

Patents Act 1977
Corrected Search Report under Section 17

Documents considered to be relevant:

Category	Relevant to claims	Identity of document and passage or figure of particular relevance
X	1-5 & 20	US 2009/232037 A1 (NOKIA SIEMENS NETWORKS) - Whole document see especially paragraphs 31 and 34
X	1-5 & 20	US 2011/053493 A1 (OKI ELECTRIC) - Whole document

Categories:

X	Document indicating lack of novelty or inventive step	A	Document indicating technological background and/or state of the art.
Y	Document indicating lack of inventive step if combined with one or more other documents of same category.	P	Document published on or after the declared priority date but before the filing date of this invention.
&	Member of the same patent family	E	Patent document published on or after, but with priority date earlier than, the filing date of this application.

Field of Search:

Search of GB, EP, WO & US patent documents classified in the following areas of the UKC^{X} :

Worldwide search of patent documents classified in the following areas of the IPC

H04W

The following online and other databases have been used in the preparation of this search report

EPODOC, WPI

International Classification:

Subclass	Subgroup	Valid From
H04W	0052/02	01/01/2009
H04W	0084/18	01/01/2009