




(22) Date de dépôt/Filing Date: 2013/01/10
(41) Mise à la disp. pub./Open to Public Insp.: 2013/08/21
(30) Priorité/Priority: 2012/02/21 (US13/356,903)

(51) Cl.Int./Int.Cl. *H04W 84/20* (2009.01),
H04L 7/00 (2006.01), *H04B 1/713* (2011.01)
(71) Demandeur/Applicant:
ELSTER SOLUTIONS, LLC, US
(72) Inventeurs/Inventors:
SHUEY, KENNETH C., US;
HEMMINGER, RODNEY C., US
(74) Agent: GOWLING LAFLEUR HENDERSON LLP

(54) Titre : PAQUETS ECHELONNABLES DANS UN SYSTEME A ETALEMENT DE SPECTRE A SAUT DE FREQUENCE
(54) Title: SCALABLE PACKETS IN A FREQUENCY HOPPING SPREAD SPECTRUM (FHSS) SYSTEM



(57) **Abrégé/Abstract:**

A wireless mesh network, method, and processor-readable storage medium for operating a network are disclosed. In particular, methods are disclosed for operating a wireless mesh network to increase the number of frequency channels used by the system

(57) **Abrégé(suite)/Abstract(continued):**

without needing to increase the preamble length. According to various embodiments, a wireless mesh network divides the ISM spectrum into two types of channels: acquisition channels and data channels. An acquisition channel is used to transmit a preamble of a communication packet. One or more data channels are used to transmit a data portion of the communication packet.

ABSTRACT

A wireless mesh network, method, and processor-readable storage medium for operating a network are disclosed. In particular, methods are disclosed for operating a wireless mesh network to increase the number of frequency channels used by the system without needing to increase the preamble length. According to various embodiments, a wireless mesh network divides the ISM spectrum into two types of channels: acquisition channels and data channels. An acquisition channel is used to transmit a preamble of a communication packet. One or more data channels are used to transmit a data portion of the communication packet.

SCALABLE PACKETS IN A FREQUENCY HOPPING SPREAD SPECTRUM (FHSS) SYSTEM

TECHNICAL BACKGROUND

[0001] The reading of electrical energy, water flow, and gas usage has historically been accomplished with human meter readers who came on-site and manually documented meter readings. Over time, this manual meter reading methodology has been enhanced with walk by or drive by reading systems that use radio communications to and from a mobile collector device in a vehicle. Recently, there has been a concerted effort to accomplish meter reading using fixed communication networks that allow data to flow from the meter to a host computer system without human intervention.

[0002] Fixed communication networks can operate using wire line or radio technology. For example, distribution line carrier systems are wire-based and use the utility lines themselves for communications. Radio technology has tended to be preferred due to higher data rates and independence from the distribution network.

[0003] Automated systems, such as Automatic Meter Reading (AMR) and Advanced Metering Infrastructure (AMI) systems, exist for collecting data from meters that measure usage of resources, such as gas, water and electricity. Such systems may employ a number of different infrastructures for collecting this meter data from the meters. For example, some automated systems obtain data from the meters using a fixed wireless network that includes, for example, a central node, *e.g.*, a collection device, in communication with a number of endpoint nodes (*e.g.*, meter reading devices (MRDs) connected to meters). At the endpoint nodes, the wireless communications circuitry may be incorporated into the meters themselves, such that each endpoint node in the wireless network comprises a meter connected to an MRD that has wireless communication circuitry that enables the MRD to transmit the meter data of the

meter to which it is connected. The wireless communication circuitry may include a transponder that is uniquely identified by a transponder serial number. The endpoint nodes may either transmit their meter data directly to the central node, or indirectly through one or more intermediate bi-directional nodes that serve as repeaters for the meter data of the transmitting node.

[0004] Some networks may employ a mesh networking architecture. In such networks, known as “mesh networks,” endpoint nodes are connected to one another through wireless communication links such that each endpoint node has a wireless communication path to the central node. One characteristic of mesh networks is that the component nodes can all connect to one another via one or more “hops.” Due to this characteristic, mesh networks can continue to operate even if a node or a connection breaks down. Accordingly, mesh networks are self-configuring and self-healing, significantly reducing installation and maintenance efforts.

[0005] Within these smart mesh networks, communications are commonly conducted in the 900 MHz Industrial, Scientific, and Medical (ISM) frequency band, which uses frequencies in the 902-928 MHz frequency range. Radio frequency (RF) technology in this range can operate without a Federal Communications Commission (FCC) license by restricting power output and by spreading the transmitted energy over a large portion of the available bandwidth. Many systems are designed by communications experts rather than by metering or utility experts. As a result, many such systems are very complex, require heavy synchronization traffic, and are difficult to troubleshoot.

[0006] While many mesh networks can perform on demand reads from the collector, the default mode of operation is *ad hoc*, in that endpoint devices establish connections with neighboring devices and initiate data transmissions directed to the collector. To eliminate contention within an *ad hoc* mesh communication system, sophisticated synchronizing schemes are required to ensure both frequency and time are kept extremely accurately. Endpoint devices need to maintain their synchronization to know exactly when to transmit and on what frequency to transmit. If tight frequency and time synchronization are not maintained, system throughput is drastically and adversely impacted.

[0007] Polled type mesh networks, such as the EnergyAxis network available from Elster Electricity, LLC, use self-coherent communications at all endpoint devices within the network and do not rely on frequency or time synchronization by the endpoint devices. This is because the endpoint devices do not need to know the time or frequency of an expected message. Each endpoint device in a polled type mesh network looks for incoming messages continually and can synchronize to both time and frequency on the fly by decoding the preamble

of an incoming message. By avoiding the need for frequency and time synchronization, the complexity of the endpoint device and of the network as a whole can be reduced. The amount of communications traffic on the network can be reduced by nearly an order of magnitude.

[0008] In a polled type mesh network, an endpoint device normally only outputs data when it is asked for information or when it has an exception message based on a change of state, alarm condition, or scheduled event. In the EnergyAxis system, several seconds are periodically set aside for exception traffic. All exception messages are decoded on the fly by endpoint meters that need to forward the data on to the collector.

[0009] Polled type mesh networks generally use polled messages and exception messages that have a fixed length. Such messages typically include a preamble time that is long enough for the endpoint device to scan all of the possible channels in order to find the channel on which the message is transmitted. After this channel is found, an algorithm is executed to accurately lock on the preamble bits and decode them until a start frame delimiter is received. The start frame delimiter marks the end of the preamble and the start of the data portion, or payload, of the message. In some cases, a different start frame delimiter is used to indicate one data rate as opposed to another. The data portion of the message, sent on the same frequency, is decoded to determine the desired function. Each new message from a given device is transmitted on a new frequency to promote the equal utilization of spectrum that is required by FCC regulations for 900 MHz ISM band operation. In addition, the new frequency complies with some pseudorandom sequence of channels.

[0010] Under current FCC regulations, a system with 25 frequency channels can transmit with up to 0.25 W power. A system with 50 frequency channels, on the other hand, can transmit with up to 1 W power. Accordingly, increasing the number of frequency channels would allow a mesh network to transmit with more power. However, as the number of channels increases, the length of the preamble required to ensure that a particular channel is found also increases.

SUMMARY OF THE DISCLOSURE

[0011] A wireless mesh network, method, and processor-readable storage medium for operating a network are disclosed. In particular, methods are disclosed for operating a wireless mesh network to increase the number of frequency channels used by the system without needing to increase the preamble length. According to various embodiments, a wireless mesh network divides the ISM spectrum into two types of channels: acquisition channels and

data channels. An acquisition channel is used to transmit a preamble of a communication packet. One or more data channels are used to transmit a data portion of the communication packet.

[0012] One embodiment is directed to a wireless network comprising a control node and a plurality of communication nodes in wireless communication with the control node. Each of the communication nodes has a wireless communication path to the control node that is either a direct path or an indirect path through one or more other communication nodes that serve as repeaters. A communication node communicates a communication packet comprising a preamble portion communicated on a dedicated acquisition channel of a frequency spectrum and a data portion communicated on at least one dedicated data channel of the frequency spectrum.

[0013] Another embodiment is directed to a method of operating a network that has a control node that communicates with a plurality of communication nodes, in which each of the communication nodes has a wireless communication path to the control node that is either a direct path or an indirect path through one or more other communication nodes that serve as repeaters. The method involves scanning a plurality of dedicated acquisition channels of a frequency spectrum to detect a preamble portion of a communication packet. The dedicated acquisition channels are ranked in descending order of respective received signal strength indicator (RSSI) values. The method also involves attempting to detect a carrier on at least some of the dedicated acquisition channels and receiving a data portion of the communication packet using at least one dedicated data channel distinct from the dedicated acquisition channels.

[0014] Yet another embodiment is directed to a processor-readable storage medium storing processor-executable instructions that, when executed by a processor, cause the processor to operate a wireless mesh network having a control node that communicates with a plurality of communication nodes, in which each of the communication nodes has a wireless communication path to the control node that is either a direct path or an indirect path through one or more other communication nodes that serve as repeaters. The instructions include instructions for scanning a plurality of dedicated acquisition channels of a frequency spectrum to detect a preamble portion of a communication packet, ranking the dedicated acquisition channels in descending order of respective received signal strength indicator (RSSI) values, attempting to detect a carrier on at least some of the dedicated acquisition channels, and receiving a data portion of the communication packet using at least one dedicated data channel distinct from the dedicated acquisition channels.

[0015] Various embodiments may realize certain advantages. For example, by using fewer acquisition channels relative to conventional systems, the preamble time can be reduced, and the overall packet time can be similarly reduced as a result. A wireless mesh

network can increase the total number of frequency channels used and thus be eligible under FCC regulations to employ more power in the transmitted signal.

[0016] Other features and advantages of the described embodiments may become apparent from the following detailed description and accompanying drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

[0017] The foregoing summary, as well as the following detailed description of various embodiments, is better understood when read in conjunction with the appended drawings. For the purpose of illustrating the invention, there are shown in the drawings exemplary embodiments of various aspects of the invention; however, the invention is not limited to the specific methods and instrumentalities disclosed. In the drawings:

[0018] Figure 1 is a diagram of an exemplary metering system;

[0019] Figure 2 expands upon the diagram of Fig. 1 and illustrates an exemplary metering system in greater detail;

[0020] Figure 3A is a block diagram illustrating an exemplary collector;

[0021] Figure 3B is a block diagram illustrating an exemplary meter;

[0022] Figure 4 is a diagram of an exemplary subnet of a wireless network for collecting data from remote devices; and

[0023] Figure 5 is a flow diagram illustrating an example method for operating a wireless mesh network.

DETAILED DESCRIPTION OF ILLUSTRATIVE EMBODIMENTS

[0024] Exemplary systems and methods for gathering meter data are described below with reference to Figures 1-5. It will be appreciated by those of ordinary skill in the art that the description given herein with respect to those figures is for exemplary purposes only and is not intended in any way to limit the scope of potential embodiments.

[0025] Generally, a plurality of meter devices, which operate to track usage of a service or commodity such as, for example, electricity, water, and gas, are operable to wirelessly communicate. One or more devices, referred to herein as “collectors,” are provided that “collect” data transmitted by the other meter devices so that it can be accessed by other computer systems. The collectors receive and compile metering data from a plurality of meter devices via wireless communications. A data collection server may communicate with the collectors to retrieve the compiled meter data.

[0026] Figure 1 provides a diagram of one exemplary metering system 110. System 110 comprises a plurality of meters 114, which are operable to sense and record

consumption or usage of a service or commodity such as, for example, electricity, water, or gas. Meters 114 may be located at customer premises such as, for example, a home or place of business. Meters 114 comprise circuitry for measuring the consumption of the service or commodity being consumed at their respective locations and for generating data reflecting the consumption, as well as other data related thereto. Meters 114 may also comprise circuitry for wirelessly transmitting data generated by the meter to a remote location. Meters 114 may further comprise circuitry for receiving data, commands or instructions wirelessly as well. Meters that are operable to both receive and transmit data may be referred to as "bi-directional" or "two-way" meters, while meters that are only capable of transmitting data may be referred to as "transmit-only" or "one-way" meters. In bi-directional meters, the circuitry for transmitting and receiving may comprise a transceiver. In an illustrative embodiment, meters 114 may be, for example, electricity meters manufactured by Elster Electricity, LLC and marketed under the tradename REX.

[0027] System 110 further comprises collectors 116. In one embodiment, collectors 116 are also meters operable to detect and record usage of a service or commodity such as, for example, electricity, water, or gas. In addition, collectors 116 are operable to send data to and receive data from meters 114. Thus, like the meters 114, the collectors 116 may comprise both circuitry for measuring the consumption of a service or commodity and for generating data reflecting the consumption and circuitry for transmitting and receiving data. In one embodiment, collector 116 and meters 114 communicate with and amongst one another using any one of several wireless techniques such as, for example, frequency hopping spread spectrum (FHSS) and direct sequence spread spectrum (DSSS).

[0028] A collector 116 and the meters 114 with which it communicates define a subnet/LAN 120 of system 110. As used herein, meters 114 and collectors 116 may be referred to as "nodes" in the subnet 120. In each subnet/LAN 120, each meter transmits data related to consumption of the commodity being metered at the meter's location. The collector 116 receives the data transmitted by each meter 114, effectively "collecting" it, and then periodically transmits the data from all of the meters in the subnet/LAN 120 to a data collection server 206. The data collection server 206 stores the data for analysis and preparation of bills, for example. The data collection server 206 may be a specially programmed general purpose computing system and may communicate with collectors 116 via a network 112. The network 112 may comprise any form of network, including a wireless network or a fixed-wire network, such as a local area network (LAN), a wide area network, the Internet, an intranet, a telephone network, such as the public switched telephone network (PSTN), a Frequency Hopping Spread Spectrum

(FHSS) radio network, a mesh network, a Wi-Fi (802.11) network, a Wi-Max (802.16) network, a land line (POTS) network, or any combination of the above.

[0029] Referring now to Figure 2, further details of the metering system 110 are shown. Typically, the system will be operated by a utility company or a company providing information technology services to a utility company. As shown, the system 110 comprises a network management server 202, a network management system (NMS) 204 and the data collection server 206 that together manage one or more subnets/LANs 120 and their constituent nodes. The NMS 204 tracks changes in network state, such as new nodes registering/unregistering with the system 110, node communication paths changing, *etc.* This information is collected for each subnet/LAN 120 and is detected and forwarded to the network management server 202 and data collection server 206.

[0030] Each of the meters 114 and collectors 116 is assigned an identifier (LAN ID) that uniquely identifies that meter or collector on its subnet/LAN 120. In this embodiment, communication between nodes (*i.e.*, the collectors and meters) and the system 110 is accomplished using the LAN ID. However, it is preferable for operators of a utility to query and communicate with the nodes using their own identifiers. To this end, a marriage file 208 may be used to correlate a utility's identifier for a node (*e.g.*, a utility serial number) with both a manufacturer serial number (*i.e.*, a serial number assigned by the manufacturer of the meter) and the LAN ID for each node in the subnet/LAN 120. In this manner, the utility can refer to the meters and collectors by the utilities identifier, while the system can employ the LAN ID for the purpose of designating particular meters during system communications.

[0031] A device configuration database 210 stores configuration information regarding the nodes. For example, in the metering system 200, the device configuration database may include data regarding time of use (TOU) switchpoints, *etc.* for the meters 114 and collectors 116 communicating in the system 110. A data collection requirements database 212 contains information regarding the data to be collected on a per node basis. For example, a utility may specify that metering data such as load profile, demand, TOU, *etc.* is to be collected from particular meter(s) 114a. Reports 214 containing information on the network configuration may be automatically generated or in accordance with a utility request.

[0032] The network management system (NMS) 204 maintains a database describing the current state of the global fixed network system (current network state 220) and a database describing the historical state of the system (historical network state 222). The current network state 220 contains data regarding current meter-to-collector assignments, *etc.* for each subnet/LAN 120. The historical network state 222 is a database from which the state of the

network at a particular point in the past can be reconstructed. The NMS 204 is responsible for, amongst other things, providing reports 214 about the state of the network. The NMS 204 may be accessed via an API 220 that is exposed to a user interface 216 and a Customer Information System (CIS) 218. Other external interfaces may also be implemented. In addition, the data collection requirements stored in the database 212 may be set via the user interface 216 or CIS 218.

[0033] The data collection server 206 collects data from the nodes (*e.g.*, collectors 116) and stores the data in a database 224. The data includes metering information, such as energy consumption and may be used for billing purposes, *etc.* by a utility provider.

[0034] The network management server 202, network management system 204 and data collection server 206 communicate with the nodes in each subnet/LAN 120 via network 110.

[0035] Figure 3A is a block diagram illustrating further details of one embodiment of a collector 116. Although certain components are designated and discussed with reference to Figure 3A, it should be appreciated that the invention is not limited to such components. In fact, various other components typically found in an electronic meter may be a part of collector 116, but have not been shown in Figure 3A for the purposes of clarity and brevity. Also, the invention may use other components to accomplish the operation of collector 116. The components that are shown and the functionality described for collector 116 are provided as examples, and are not meant to be exclusive of other components or other functionality.

[0036] As shown in Figure 3A, collector 116 may comprise metering circuitry 304 that performs measurement of consumption of a service or commodity and a processor 305 that controls the overall operation of the metering functions of the collector 116. The collector 116 may further comprise a display 310 for displaying information such as measured quantities and meter status and a memory 312 for storing data. The collector 116 further comprises wireless LAN communications circuitry 306 for communicating wirelessly with the meters 114 in a subnet/LAN and a network interface 308 for communication over the network 112.

[0037] In one embodiment, the metering circuitry 304, processor 305, display 310 and memory 312 are implemented using an A3 ALPHA meter available from Elster Electricity, Inc. In that embodiment, the wireless LAN communications circuitry 306 may be implemented by a LAN Option Board (*e.g.*, a 900 MHz two-way radio) installed within the A3 ALPHA meter, and the network interface 308 may be implemented by a WAN Option Board (*e.g.*, a telephone modem) also installed within the A3 ALPHA meter. In this embodiment, the WAN Option

Board 308 routes messages from network 112 (via interface port 302) to either the meter processor 305 or the LAN Option Board 306. LAN Option Board 306 may use a transceiver (not shown), for example a 900 MHz radio, to communicate data to meters 114. Also, LAN Option Board 306 may have sufficient memory to store data received from meters 114. This data may include, but is not limited to the following: current billing data (*e.g.*, the present values stored and displayed by meters 114), previous billing period data, previous season data, and load profile data.

[0038] LAN Option Board 306 may be capable of synchronizing its time to a real time clock (not shown) in A3 ALPHA meter, thereby synchronizing the LAN reference time to the time in the meter. The processing necessary to carry out the communication functionality and the collection and storage of metering data of the collector 116 may be handled by the processor 305 and/or additional processors (not shown) in the LAN Option Board 306 and the WAN Option Board 308.

[0039] The responsibility of a collector 116 is wide and varied. Generally, collector 116 is responsible for managing, processing and routing data communicated between the collector and network 112 and between the collector and meters 114. Collector 116 may continually or intermittently read the current data from meters 114 and store the data in a database (not shown) in collector 116. Such current data may include but is not limited to the total kWh usage, the Time-Of-Use (TOU) kWh usage, peak kW demand, and other energy consumption measurements and status information. Collector 116 also may read and store previous billing and previous season data from meters 114 and store the data in the database in collector 116. The database may be implemented as one or more tables of data within the collector 116.

[0040] Figure 3B is a block diagram of an exemplary embodiment of a meter 114 that may operate in the system 110 of Figures 1 and 2. As shown, the meter 114 comprises metering circuitry 304' for measuring the amount of a service or commodity that is consumed, a processor 305' that controls the overall functions of the meter, a display 310' for displaying meter data and status information, and a memory 312' for storing data and program instructions. The meter 114 further comprises wireless communications circuitry 306' for transmitting and receiving data to/from other meters 114 or a collector 116.

[0041] Referring again to Figure 1, in the exemplary embodiment shown, a collector 116 directly communicates with only a subset of the plurality of meters 114 in its particular subnet/LAN. Meters 114 with which collector 116 directly communicates may be referred to as "level one" meters 114a. The level one meters 114a are said to be one "hop" from

the collector 116. Communications between collector 116 and meters 114 other than level one meters 114a are relayed through the level one meters 114a. Thus, the level one meters 114a operate as repeaters for communications between collector 116 and meters 114 located further away in subnet 120.

[0042] Each level one meter 114a typically will only be in range to directly communicate with only a subset of the remaining meters 114 in the subnet 120. The meters 114 with which the level one meters 114a directly communicate may be referred to as level two meters 114b. Level two meters 114b are one “hop” from level one meters 114a, and therefore two “hops” from collector 116. Level two meters 114b operate as repeaters for communications between the level one meters 114a and meters 114 located further away from collector 116 in the subnet 120.

[0043] While only three levels of meters are shown (collector 116, first level 114a, second level 114b) in Figure 1, a subnet 120 may comprise any number of levels of meters 114. For example, a subnet 120 may comprise one level of meters but might also comprise eight or more levels of meters 114. In an embodiment wherein a subnet comprises eight levels of meters 114, as many as 1024 meters might be registered with a single collector 116.

[0044] As mentioned above, each meter 114 and collector 116 that is installed in the system 110 has a unique identifier (LAN ID) stored thereon that uniquely identifies the device from all other devices in the system 110. Additionally, meters 114 operating in a subnet 120 comprise information including the following: data identifying the collector with which the meter is registered; the level in the subnet at which the meter is located; the repeater meter at the prior level with which the meter communicates to send and receive data to/from the collector; an identifier indicating whether the meter is a repeater for other nodes in the subnet; and if the meter operates as a repeater, the identifier that uniquely identifies the repeater within the particular subnet, and the number of meters for which it is a repeater. Collectors 116 have stored thereon all of this same data for all meters 114 that are registered therewith. Thus, collector 116 comprises data identifying all nodes registered therewith as well as data identifying the registered path by which data is communicated from the collector to each node. Each meter 114 therefore has a designated communications path to the collector that is either a direct path (*e.g.*, all level one nodes) or an indirect path through one or more intermediate nodes that serve as repeaters.

[0045] Information is transmitted in this embodiment in the form of packets. For most network tasks such as, for example, reading meter data, collector 116 communicates with meters 114 in the subnet 120 using point-to-point transmissions. For example, a message or

instruction from collector 116 is routed through the designated set of repeaters to the desired meter 114. Similarly, a meter 114 communicates with collector 116 through the same set of repeaters, but in reverse.

[0046] In some instances, however, collector 116 may need to quickly communicate information to all meters 114 located in its subnet 120. Accordingly, collector 116 may issue a broadcast message that is meant to reach all nodes in the subnet 120. The broadcast message may be referred to as a “flood broadcast message.” A flood broadcast originates at collector 116 and propagates through the entire subnet 120 one level at a time. For example, collector 116 may transmit a flood broadcast to all first level meters 114a. The first level meters 114a that receive the message pick a random time slot and retransmit the broadcast message to second level meters 114b. Any second level meter 114b can accept the broadcast, thereby providing better coverage from the collector out to the end point meters. Similarly, the second level meters 114b that receive the broadcast message pick a random time slot and communicate the broadcast message to third level meters. This process continues out until the end nodes of the subnet. Thus, a broadcast message gradually propagates outward from the collector to the nodes of the subnet 120.

[0047] The flood broadcast packet header contains information to prevent nodes from repeating the flood broadcast packet more than once per level. For example, within a flood broadcast message, a field might exist that indicates to meters/nodes which receive the message, the level of the subnet the message is located; only nodes at that particular level may re-broadcast the message to the next level. If the collector broadcasts a flood message with a level of 1, only level 1 nodes may respond. Prior to re-broadcasting the flood message, the level 1 nodes increment the field to 2 so that only level 2 nodes respond to the broadcast. Information within the flood broadcast packet header ensures that a flood broadcast will eventually die out.

[0048] Generally, a collector 116 issues a flood broadcast several times, *e.g.* five times, successively to increase the probability that all meters in the subnet 120 receive the broadcast. A delay is introduced before each new broadcast to allow the previous broadcast packet time to propagate through all levels of the subnet.

[0049] Meters 114 may have a clock formed therein. However, meters 114 often undergo power interruptions that can interfere with the operation of any clock therein. Accordingly, the clocks internal to meters 114 cannot be relied upon to provide an accurate time reading. Having the correct time is necessary, however, when time of use metering is being employed. Indeed, in an embodiment, time of use schedule data may also be comprised in the same broadcast message as the time. Accordingly, collector 116 periodically flood broadcasts

the real time to meters 114 in subnet 120. Meters 114 use the time broadcasts to stay synchronized with the rest of the subnet 120. In an illustrative embodiment, collector 116 broadcasts the time every 15 minutes. The broadcasts may be made near the middle of 15 minute clock boundaries that are used in performing load profiling and time of use (TOU) schedules so as to minimize time changes near these boundaries. Maintaining time synchronization is important to the proper operation of the subnet 120. Accordingly, lower priority tasks performed by collector 116 may be delayed while the time broadcasts are performed.

[0050] In an illustrative embodiment, the flood broadcasts transmitting time data may be repeated, for example, five times, so as to increase the probability that all nodes receive the time. Furthermore, where time of use schedule data is communicated in the same transmission as the timing data, the subsequent time transmissions allow a different piece of the time of use schedule to be transmitted to the nodes.

[0051] Exception messages are used in subnet 120 to transmit unexpected events that occur at meters 114 to collector 116. In an embodiment, the first 4 seconds of every 32-second period are allocated as an exception window for meters 114 to transmit exception messages. Meters 114 transmit their exception messages early enough in the exception window so the message has time to propagate to collector 116 before the end of the exception window. Collector 116 may process the exceptions after the 4-second exception window. Generally, a collector 116 acknowledges exception messages, and collector 116 waits until the end of the exception window to send this acknowledgement.

[0052] In an illustrative embodiment, exception messages are configured as one of three different types of exception messages: local exceptions, which are handled directly by the collector 116 without intervention from data collection server 206; an immediate exception, which is generally relayed to data collection server 206 under an expedited schedule; and a daily exception, which is communicated to the communication server 122 on a regular schedule.

[0053] Exceptions are processed as follows. When an exception is received at collector 116, the collector 116 identifies the type of exception that has been received. If a local exception has been received, collector 116 takes an action to remedy the problem. For example, when collector 116 receives an exception requesting a "node scan request" such as discussed below, collector 116 transmits a command to initiate a scan procedure to the meter 114 from which the exception was received.

[0054] If an immediate exception type has been received, collector 116 makes a record of the exception. An immediate exception might identify, for example, that there has

been a power outage. Collector 116 may log the receipt of the exception in one or more tables or files. In an illustrative example, a record of receipt of an immediate exception is made in a table referred to as the "Immediate Exception Log Table." Collector 116 then waits a set period of time before taking further action with respect to the immediate exception. For example, collector 116 may wait 64 seconds. This delay period allows the exception to be corrected before communicating the exception to the data collection server 206. For example, where a power outage was the cause of the immediate exception, collector 116 may wait a set period of time to allow for receipt of a message indicating the power outage has been corrected.

[0055] If the exception has not been corrected, collector 116 communicates the immediate exception to data collection server 206. For example, collector 116 may initiate a dial-up connection with data collection server 206 and download the exception data. After reporting an immediate exception to data collection server 206, collector 116 may delay reporting any additional immediate exceptions for a period of time such as ten minutes. This is to avoid reporting exceptions from other meters 114 that relate to, or have the same cause as, the exception that was just reported.

[0056] If a daily exception was received, the exception is recorded in a file or a database table. Generally, daily exceptions are occurrences in the subnet 120 that need to be reported to data collection server 206, but are not so urgent that they need to be communicated immediately. For example, when collector 116 registers a new meter 114 in subnet 120, collector 116 records a daily exception identifying that the registration has taken place. In an illustrative embodiment, the exception is recorded in a database table referred to as the "Daily Exception Log Table." Collector 116 communicates the daily exceptions to data collection server 206. Generally, collector 116 communicates the daily exceptions once every 24 hours.

[0057] In the present embodiment, a collector assigns designated communications paths to meters with bi-directional communication capability, and may change the communication paths for previously registered meters if conditions warrant. For example, when a collector 116 is initially brought into system 110, it needs to identify and register meters in its subnet 120. A "node scan" refers to a process of communication between a collector 116 and meters 114 whereby the collector may identify and register new nodes in a subnet 120 and allow previously registered nodes to switch paths. A collector 116 can implement a node scan on the entire subnet, referred to as a "full node scan," or a node scan can be performed on specially identified nodes, referred to as a "node scan retry."

[0058] A full node scan may be performed, for example, when a collector is first installed. The collector 116 must identify and register nodes from which it will collect usage

data. The collector 116 initiates a node scan by broadcasting a request, which may be referred to as a Node Scan Procedure request. Generally, the Node Scan Procedure request directs that all unregistered meters 114 or nodes that receive the request respond to the collector 116. The request may comprise information such as the unique address of the collector that initiated the procedure. The signal by which collector 116 transmits this request may have limited strength and therefore is detected only at meters 114 that are in proximity of collector 116. Meters 114 that receive the Node Scan Procedure request respond by transmitting their unique identifier as well as other data.

[0059] For each meter from which the collector receives a response to the Node Scan Procedure request, the collector tries to qualify the communications path to that meter before registering the meter with the collector. That is, before registering a meter, the collector 116 attempts to determine whether data communications with the meter will be sufficiently reliable. In one embodiment, the collector 116 determines whether the communication path to a responding meter is sufficiently reliable by comparing a Received Signal Strength Indication (RSSI) value (*i.e.*, a measurement of the received radio signal strength) measured with respect to the received response from the meter to a selected threshold value. For example, the threshold value may be -60 dBm. RSSI values above this threshold would be deemed sufficiently reliable. In another embodiment, qualification is performed by transmitting a predetermined number of additional packets to the meter, such as ten packets, and counting the number of acknowledgements received back from the meter. If the number of acknowledgments received is greater than or equal to a selected threshold (*e.g.*, 8 out of 10), then the path is considered to be reliable. In other embodiments, a combination of the two qualification techniques may be employed.

[0060] If the qualification threshold is not met, the collector 116 may add an entry for the meter to a "Straggler Table." The entry includes the meter's LAN ID, its qualification score (*e.g.*, 5 out of 10; or its RSSI value), its level (in this case level one) and the unique ID of its parent (in this case the collector's ID).

[0061] If the qualification threshold is met or exceeded, the collector 116 registers the node. Registering a meter 114 comprises updating a list of the registered nodes at collector 116. For example, the list may be updated to identify the meter's system-wide unique identifier and the communication path to the node. Collector 116 also records the meter's level in the subnet (*i.e.* whether the meter is a level one node, level two node, *etc.*), whether the node operates as a repeater, and if so, the number of meters for which it operates as a repeater. The registration process further comprises transmitting registration information to the meter 114. For

example, collector 116 forwards to meter 114 an indication that it is registered, the unique identifier of the collector with which it is registered, the level the meter exists at in the subnet, and the unique identifier of its parent meter that will server as a repeater for messages the meter may send to the collector. In the case of a level one node, the parent is the collector itself. The meter stores this data and begins to operate as part of the subnet by responding to commands from its collector 116.

[0062] Qualification and registration continues for each meter that responds to the collector's initial Node Scan Procedure request. The collector 116 may rebroadcast the Node Scan Procedure additional times so as to insure that all meters 114 that may receive the Node Scan Procedure have an opportunity for their response to be received and the meter qualified as a level one node at collector 116.

[0063] The node scan process then continues by performing a similar process as that described above at each of the now registered level one nodes. This process results in the identification and registration of level two nodes. After the level two nodes are identified, a similar node scan process is performed at the level two nodes to identify level three nodes, and so on.

[0064] Specifically, to identify and register meters that will become level two meters, for each level one meter, in succession, the collector 116 transmits a command to the level one meter, which may be referred to as an "Initiate Node Scan Procedure" command. This command instructs the level one meter to perform its own node scan process. The request comprises several data items that the receiving meter may use in completing the node scan. For example, the request may comprise the number of timeslots available for responding nodes, the unique address of the collector that initiated the request, and a measure of the reliability of the communications between the target node and the collector. As described below, the measure of reliability may be employed during a process for identifying more reliable paths for previously registered nodes.

[0065] The meter that receives the Initiate Node Scan Response request responds by performing a node scan process similar to that described above. More specifically, the meter broadcasts a request to which all unregistered nodes may respond. The request comprises the number of timeslots available for responding nodes (which is used to set the period for the node to wait for responses), the unique address of the collector that initiated the node scan procedure, a measure of the reliability of the communications between the sending node and the collector (which may be used in the process of determining whether a meter's path may be switched as described below), the level within the subnet of the node sending the request, and an RSSI

threshold (which may also be used in the process of determining whether a registered meter's path may be switched). The meter issuing the node scan request then waits for and receives responses from unregistered nodes. For each response, the meter stores in memory the unique identifier of the responding meter. This information is then transmitted to the collector.

[0066] For each unregistered meter that responded to the node scan issued by the level one meter, the collector attempts again to determine the reliability of the communication path to that meter. In one embodiment, the collector sends a "Qualify Nodes Procedure" command to the level one node which instructs the level one node to transmit a predetermined number of additional packets to the potential level two node and to record the number of acknowledgements received back from the potential level two node. This qualification score (*e.g.*, 8 out of 10) is then transmitted back to the collector, which again compares the score to a qualification threshold. In other embodiments, other measures of the communications reliability may be provided, such as an RSSI value.

[0067] If the qualification threshold is not met, then the collector adds an entry for the node in the Straggler Table, as discussed above. However, if there already is an entry in the Straggler Table for the node, the collector will update that entry only if the qualification score for this node scan procedure is better than the recorded qualification score from the prior node scan that resulted in an entry for the node.

[0068] If the qualification threshold is met or exceeded, the collector 116 registers the node. Again, registering a meter 114 at level two comprises updating a list of the registered nodes at collector 116. For example, the list may be updated to identify the meter's unique identifier and the level of the meter in the subnet. Additionally, the collector's 116 registration information is updated to reflect that the meter 114 from which the scan process was initiated is identified as a repeater (or parent) for the newly registered node. The registration process further comprises transmitting information to the newly registered meter as well as the meter that will serve as a repeater for the newly added node. For example, the node that issued the node scan response request is updated to identify that it operates as a repeater and, if it was previously registered as a repeater, increments a data item identifying the number of nodes for which it serves as a repeater. Thereafter, collector 116 forwards to the newly registered meter an indication that it is registered, an identification of the collector 116 with which it is registered, the level the meter exists at in the subnet, and the unique identifier of the node that will serve as its parent, or repeater, when it communicates with the collector 116.

[0069] The collector then performs the same qualification procedure for each other potential level two node that responded to the level one node's node scan request. Once

that process is completed for the first level one node, the collector initiates the same procedure at each other level one node until the process of qualifying and registering level two nodes has been completed at each level one node. Once the node scan procedure has been performed by each level one node, resulting in a number of level two nodes being registered with the collector, the collector will then send the Initiate Node Scan Response command to each level two node, in turn. Each level two node will then perform the same node scan procedure as performed by the level one nodes, potentially resulting in the registration of a number of level three nodes. The process is then performed at each successive node, until a maximum number of levels is reached (*e.g.*, seven levels) or no unregistered nodes are left in the subnet.

[0070] It will be appreciated that in the present embodiment, during the qualification process for a given node at a given level, the collector qualifies the last “hop” only. For example, if an unregistered node responds to a node scan request from a level four node, and therefore, becomes a potential level five node, the qualification score for that node is based on the reliability of communications between the level four node and the potential level five node (*i.e.*, packets transmitted by the level four node versus acknowledgments received from the potential level five node), not based on any measure of the reliability of the communications over the full path from the collector to the potential level five node. In other embodiments, of course, the qualification score could be based on the full communication path.

[0071] At some point, each meter will have an established communication path to the collector which will be either a direct path (*i.e.*, level one nodes) or an indirect path through one or more intermediate nodes that serve as repeaters. If during operation of the network, a meter registered in this manner fails to perform adequately, it may be assigned a different path or possibly to a different collector as described below.

[0072] As previously mentioned, a full node scan may be performed when a collector 116 is first introduced to a network. At the conclusion of the full node scan, a collector 116 will have registered a set of meters 114 with which it communicates and reads metering data. Full node scans might be periodically performed by an installed collector to identify new meters 114 that have been brought on-line since the last node scan and to allow registered meters to switch to a different path.

[0073] In addition to the full node scan, collector 116 may also perform a process of scanning specific meters 114 in the subnet 120, which is referred to as a “node scan retry.” For example, collector 116 may issue a specific request to a meter 114 to perform a node scan outside of a full node scan when on a previous attempt to scan the node, the collector 116 was unable to confirm that the particular meter 114 received the node scan request. Also, a collector

116 may request a node scan retry of a meter 114 when during the course of a full node scan the collector 116 was unable to read the node scan data from the meter 114. Similarly, a node scan retry will be performed when an exception procedure requesting an immediate node scan is received from a meter 114.

[0074] The system 110 also automatically reconfigures to accommodate a new meter 114 that may be added. More particularly, the system identifies that the new meter has begun operating and identifies a path to a collector 116 that will become responsible for collecting the metering data. Specifically, the new meter will broadcast an indication that it is unregistered. In one embodiment, this broadcast might be, for example, embedded in, or relayed as part of a request for an update of the real time as described above. The broadcast will be received at one of the registered meters 114 in proximity to the meter that is attempting to register. The registered meter 114 forwards the time to the meter that is attempting to register. The registered node also transmits an exception request to its collector 116 requesting that the collector 116 implement a node scan, which presumably will locate and register the new meter. The collector 116 then transmits a request that the registered node perform a node scan. The registered node will perform the node scan, during which it requests that all unregistered nodes respond. Presumably, the newly added, unregistered meter will respond to the node scan. When it does, the collector will then attempt to qualify and then register the new node in the same manner as described above.

[0075] Once a communication path between the collector and a meter is established, the meter can begin transmitting its meter data to the collector and the collector can transmit data and instructions to the meter. As mentioned above, data is transmitted in packets. "Outbound" packets are packets transmitted from the collector to a meter at a given level. In one embodiment, outbound packets contain the following fields, but other fields may also be included:

Length – the length of the packet;

SrcAddr – source address – in this case, the ID of the collector;

DestAddr – the LAN ID of the meter to which the packet addressed;

RptPath – the communication path to the destination meter (*i.e.*, the list of identifiers of each repeater in the path from the collector to the destination node); and

Data – the payload of the packet.

The packet may also include integrity check information (*e.g.*, CRC), a pad to fill-out unused portions of the packet and other control information. When the packet is transmitted from the collector, it will only be forwarded on to the destination meter by those repeater meters whose

identifiers appear in the RptPath field. Other meters that may receive the packet, but that are not listed in the path identified in the RptPath field will not repeat the packet.

[0076] “Inbound” packets are packets transmitted from a meter at a given level to the collector. In one embodiment, inbound packets contain the following fields, but other fields may also be included:

Length – the length of the packet;

SrcAddr – source address – the address of the meter that initiated the packet;

DestAddr – the ID of the collector to which the packet is to be transmitted;

RptAddr – the ID of the parent node that serves as the next repeater for the sending node;

Data – the payload of the packet;

Because each meter knows the identifier of its parent node (*i.e.*, the node in the next lower level that serves as a repeater for the present node), an inbound packet need only identify who is the next parent. When a node receives an inbound packet, it checks to see if the RptAddr matches its own identifier. If not, it discards the packet. If so, it knows that it is supposed to forward the packet on toward the collector. The node will then replace the RptAddr field with the identifier of its own parent and will then transmit the packet so that its parent will receive it. This process will continue through each repeater at each successive level until the packet reaches the collector.

[0077] For example, suppose a meter at level three initiates transmission of a packet destined for its collector. The level three node will insert in the RptAddr field of the inbound packet the identifier of the level two node that serves as a repeater for the level three node. The level three node will then transmit the packet. Several level two nodes may receive the packet, but only the level two node having an identifier that matches the identifier in the RptAddr field of the packet will acknowledge it. The other will discard it. When the level two node with the matching identifier receives the packet, it will replace the RptAddr field of the packet with the identifier of the level one packet that serves as a repeater for that level two packet, and the level two packet will then transmit the packet. This time, the level one node having the identifier that matches the RptAddr field will receive the packet. The level one node will insert the identifier of the collector in the RptAddr field and will transmit the packet. The collector will then receive the packet to complete the transmission.

[0078] A collector 116 periodically retrieves meter data from the meters that are registered with it. For example, meter data may be retrieved from a meter every 4 hours. Where there is a problem with reading the meter data on the regularly scheduled interval, the collector will try to read the data again before the next regularly scheduled interval. Nevertheless, there

may be instances wherein the collector 116 is unable to read metering data from a particular meter 114 for a prolonged period of time. The meters 114 store an indication of when they are read by their collector 116 and keep track of the time since their data has last been collected by the collector 116. If the length of time since the last reading exceeds a defined threshold, such as for example, 18 hours, presumably a problem has arisen in the communication path between the particular meter 114 and the collector 116. Accordingly, the meter 114 changes its status to that of an unregistered meter and attempts to locate a new path to a collector 116 via the process described above for a new node. Thus, the exemplary system is operable to reconfigure itself to address inadequacies in the system.

[0079] In some instances, while a collector 116 may be able to retrieve data from a registered meter 114 occasionally, the level of success in reading the meter may be inadequate. For example, if a collector 116 attempts to read meter data from a meter 114 every 4 hours but is able to read the data, for example, only 70 percent of the time or less, it may be desirable to find a more reliable path for reading the data from that particular meter. Where the frequency of reading data from a meter 114 falls below a desired success level, the collector 116 transmits a message to the meter 114 to respond to node scans going forward. The meter 114 remains registered but will respond to node scans in the same manner as an unregistered node as described above. In other embodiments, all registered meters may be permitted to respond to node scans, but a meter will only respond to a node scan if the path to the collector through the meter that issued the node scan is shorter (*i.e.*, less hops) than the meter's current path to the collector. A lesser number of hops is assumed to provide a more reliable communication path than a longer path. A node scan request always identifies the level of the node that transmits the request, and using that information, an already registered node that is permitted to respond to node scans can determine if a potential new path to the collector through the node that issued the node scan is shorter than the node's current path to the collector.

[0080] If an already registered meter 114 responds to a node scan procedure, the collector 116 recognizes the response as originating from a registered meter but that by re-registering the meter with the node that issued the node scan, the collector may be able to switch the meter to a new, more reliable path. The collector 116 may verify that the RSSI value of the node scan response exceeds an established threshold. If it does not, the potential new path will be rejected. However, if the RSSI threshold is met, the collector 116 will request that the node that issued the node scan perform the qualification process described above (*i.e.*, send a predetermined number of packets to the node and count the number of acknowledgements received). If the resulting qualification score satisfies a threshold, then the collector will register

the node with the new path. The registration process comprises updating the collector 116 and meter 114 with data identifying the new repeater (*i.e.* the node that issued the node scan) with which the updated node will now communicate. Additionally, if the repeater has not previously performed the operation of a repeater, the repeater would need to be updated to identify that it is a repeater. Likewise, the repeater with which the meter previously communicated is updated to identify that it is no longer a repeater for the particular meter 114. In other embodiments, the threshold determination with respect to the RSSI value may be omitted. In such embodiments, only the qualification of the last “hop” (*i.e.*, sending a predetermined number of packets to the node and counting the number of acknowledgements received) will be performed to determine whether to accept or reject the new path.

[0081] In some instances, a more reliable communication path for a meter may exist through a collector other than that with which the meter is registered. A meter may automatically recognize the existence of the more reliable communication path, switch collectors, and notify the previous collector that the change has taken place. The process of switching the registration of a meter from a first collector to a second collector begins when a registered meter 114 receives a node scan request from a collector 116 other than the one with which the meter is presently registered. Typically, a registered meter 114 does not respond to node scan requests. However, if the request is likely to result in a more reliable transmission path, even a registered meter may respond. Accordingly, the meter determines if the new collector offers a potentially more reliable transmission path. For example, the meter 114 may determine if the path to the potential new collector 116 comprises fewer hops than the path to the collector with which the meter is registered. If not, the path may not be more reliable and the meter 114 will not respond to the node scan. The meter 114 might also determine if the RSSI of the node scan packet exceeds an RSSI threshold identified in the node scan information. If so, the new collector may offer a more reliable transmission path for meter data. If not, the transmission path may not be acceptable and the meter may not respond. Additionally, if the reliability of communication between the potential new collector and the repeater that would service the meter meets a threshold established when the repeater was registered with its existing collector, the communication path to the new collector may be more reliable. If the reliability does not exceed this threshold, however, the meter 114 does not respond to the node scan.

[0082] If it is determined that the path to the new collector may be better than the path to its existing collector, the meter 114 responds to the node scan. Included in the response is information regarding any nodes for which the particular meter may operate as a repeater. For

example, the response might identify the number of nodes for which the meter serves as a repeater.

[0083] The collector 116 then determines if it has the capacity to service the meter and any meters for which it operates as a repeater. If not, the collector 116 does not respond to the meter that is attempting to change collectors. If, however, the collector 116 determines that it has capacity to service the meter 114, the collector 116 stores registration information about the meter 114. The collector 116 then transmits a registration command to meter 114. The meter 114 updates its registration data to identify that it is now registered with the new collector. The collector 116 then communicates instructions to the meter 114 to initiate a node scan request. Nodes that are unregistered, or that had previously used meter 114 as a repeater respond to the request to identify themselves to collector 116. The collector registers these nodes as is described above in connection with registering new meters/nodes.

[0084] Under some circumstances it may be necessary to change a collector. For example, a collector may be malfunctioning and need to be taken off-line. Accordingly, a new communication path must be provided for collecting meter data from the meters serviced by the particular collector. The process of replacing a collector is performed by broadcasting a message to unregister, usually from a replacement collector, to all of the meters that are registered with the collector that is being removed from service. In one embodiment, registered meters may be programmed to only respond to commands from the collector with which they are registered. Accordingly, the command to unregister may comprise the unique identifier of the collector that is being replaced. In response to the command to unregister, the meters begin to operate as unregistered meters and respond to node scan requests. To allow the unregister command to propagate through the subnet, when a node receives the command it will not unregister immediately, but rather remain registered for a defined period, which may be referred to as the "Time to Live". During this time to live period, the nodes continue to respond to application layer and immediate retries allowing the unregistration command to propagate to all nodes in the subnet. Ultimately, the meters register with the replacement collector using the procedure described above.

[0085] One of collector's 116 main responsibilities within subnet 120 is to retrieve metering data from meters 114. In one embodiment, collector 116 has as a goal to obtain at least one successful read of the metering data per day from each node in its subnet. Collector 116 attempts to retrieve the data from all nodes in its subnet 120 at a configurable periodicity. For example, collector 116 may be configured to attempt to retrieve metering data from meters 114 in its subnet 120 once every 4 hours. In greater detail, in one embodiment, the

data collection process begins with the collector 116 identifying one of the meters 114 in its subnet 120. For example, collector 116 may review a list of registered nodes and identify one for reading. The collector 116 then communicates a command to the particular meter 114 that it forward its metering data to the collector 116. If the meter reading is successful and the data is received at collector 116, the collector 116 determines if there are other meters that have not been read during the present reading session. If so, processing continues. However, if all of the meters 114 in subnet 120 have been read, the collector waits a defined length of time, such as, for example, 4 hours, before attempting another read.

[0086] If during a read of a particular meter, the meter data is not received at collector 116, the collector 116 begins a retry procedure wherein it attempts to retry the data read from the particular meter. Collector 116 continues to attempt to read the data from the node until either the data is read or the next subnet reading takes place. In an embodiment, collector 116 attempts to read the data every 60 minutes. Thus, wherein a subnet reading is taken every 4 hours, collector 116 may issue three retries between subnet readings.

[0087] Meters 114 are often two-way meters – *i.e.* they are operable to both receive and transmit data. However, one-way meters that are operable only to transmit and not receive data may also be deployed. Figure 4 is a block diagram illustrating a subnet 401 that includes a number of one-way meters 451-456. As shown, meters 114a-k are two-way devices. In this example, the two-way meters 114a-k operate in the exemplary manner described above, such that each meter has a communication path to the collector 116 that is either a direct path (*e.g.*, meters 114a and 114b have a direct path to the collector 116) or an indirect path through one or more intermediate meters that serve as repeaters. For example, meter 114h has a path to the collector through, in sequence, intermediate meters 114d and 114b. In this example embodiment, when a one-way meter (*e.g.*, meter 451) broadcasts its usage data, the data may be received at one or more two-way meters that are in proximity to the one-way meter (*e.g.*, two-way meters 114f and 114g). In one embodiment, the data from the one-way meter is stored in each two-way meter that receives it, and the data is designated in those two-way meters as having been received from the one-way meter. At some point, the data from the one-way meter is communicated, by each two-way meter that received it, to the collector 116. For example, when the collector reads the two-way meter data, it recognizes the existence of meter data from the one-way meter and reads it as well. After the data from the one-way meter has been read, it is removed from memory.

[0088] While the collection of data from one-way meters by the collector has been described above in the context of a network of two-way meters 114 that operate in the

manner described in connection with the embodiments described above, it is understood that the present invention is not limited to the particular form of network established and utilized by the meters 114 to transmit data to the collector. Rather, the present invention may be used in the context of any network topology in which a plurality of two-way communication nodes are capable of transmitting data and of having that data propagated through the network of nodes to the collector.

[0089] According to various example embodiments disclosed herein, the ISM spectrum is separated into a set of acquisition channels and a set of data channels. A communication packet is transmitted using one of the acquisition channels for the preamble. One or more of the data channels is used for the data portion of the message. Because fewer acquisition channels are used as compared with conventional systems, the preamble time can be reduced. As a result, the overall packet time can also be reduced. The number of acquisition channels can be selected to realize a compromise between a sufficient number of channels to minimize contention on the one hand, and a low enough number of channels to keep a short preamble length, on the other hand. Accordingly, the embodiments disclosed herein improve on conventional systems by reducing overall packet time, while still meeting FCC equal channel requirements.

[0090] Figure 5 illustrates an example method 500 for operating a wireless network according to an embodiment. At a step 502, a receiving device scans a subset of the frequency channels in the ISM band for a signal. In particular, a subset of channels that have been designated *a priori* as dedicated acquisition channels is scanned. This approach differs from conventional approaches in which all channels in the ISM band are scanned, *i.e.*, all channels effectively serve as both acquisition channels and data channels. Scanning only a subset of the channels rather than all of the channels allows the preamble length, and thus the overall packet length, to be reduced.

[0091] Next, at a step 504, the acquisition channels are ranked in decreasing order of signal strength, as indicated by their respective received signal strength indicator (RSSI) values. The receiving device then attempts to detect a carrier on some or all of the acquisition channels, in decreasing order of signal strength, at a step 506. In some conventional approaches, by comparison, the device would attempt to detect a carrier on the twelve channels (out of twenty-five total channels) having the highest RSSI values. Accordingly, messages transmitted in conventional systems include a preamble that is long enough for the endpoint receiver to scan all possible channels to find the channel on which the message is transmitted. Searching for a

carrier on a reduced number of channels relative to conventional approaches can reduce the time involved in searching for a carrier. The preamble length can be reduced as a result.

[0092] After the carrier is detected, the receiving device receives the data portion of the message on one or more of the dedicated data channels, which are the frequency channels of the ISM band that are not designated as acquisition channels, at a step 508. The numbers of acquisition channels and data channels may be selected to achieve a balance between reducing contention and reducing preamble length, as demonstrated by the examples disclosed below.

[0093] The data in Table 1, below, shows how a variable preamble time can be scaled to match different fixed packet lengths to correctly average channel utilization across the ISM frequency band.

| | | |
|---|------|-----|
| Data packet length (msec) | 20 | 100 |
| Total number of channels | 50 | 50 |
| Acquisition channels | 4 | 4 |
| Data channels | 46 | 46 |
| Preamble length for balanced channel usage (msec) | 1.74 | 8.7 |

Table 1

[0094] In the first example in Table 1, for a 20 msec data packet, a 1.74 msec preamble can be used in a 50 channel system that has four acquisition channels and 46 data channels. After a total of 184 (4 x 26) data packets from a given device, the average time occupancy – that is, the amount of time spent transmitting on each channel – on all 50 channels will be substantially equal. In the second example of Table 1, the fixed data packet length is increased to 100 msec. To accommodate this increase, the preamble length is also increased, to 8.7 msec. With these parameters, after 184 data packets from a given device, all 50 channels will have been substantially used equally, on average.

[0095] The following Table 2 shows two example embodiments that employ six acquisition channels and 44 data channels in a 50 channel system, as compared with four acquisition channels and 46 data channels as shown in Table 1. It can be seen that, to accommodate the increased number of acquisition channels, the preamble length also needs to increase. In other words, according to an aspect of the disclosed subject matter, the preamble length scales with the number of acquisition channels in the system. With these parameters, the FCC equal channel utilization requirement is achieved after 264 data packets, as compared with 184 data packets in the examples in Table 1. Thus, increasing the number of acquisition

channels increases both the required preamble length and the number of data packets that are transmitted before the FCC equal channel utilization requirement is achieved.

| | | |
|---|-----|------|
| Data packet length (msec) | 20 | 100 |
| Total number of channels | 50 | 50 |
| Acquisition channels | 6 | 6 |
| Data channels | 44 | 44 |
| Preamble length for balanced channel usage (msec) | 2.7 | 13.6 |

Table 2

[0096] Table 3, below, illustrates another example embodiment of the disclosed scaling concept that is applicable to a 50 channel FHSS system. In this embodiment, the data packet is substantially longer – 200 msec. Accordingly, the data packet portion of the message is transmitted on two different data channels, as compared with a single data channel as shown in the examples of Tables 1 and 2 above. Each message comprises a preamble on one of the four dedicated acquisition channels. The preamble is followed by 100 msec of data on one data channel and 100 msec of data on another data channel. Using this technique of transmitting data packets on multiple data channels decreases the number of data packets that are transmitted before the FCC equal channel utilization requirement is achieved. Specifically, all 50 channels would be used substantially equally, on average, after 92 data packets from any transceiver.

| | |
|---|-----------------------------|
| Data packet length (msec) | 200 |
| Total number of channels | 50 |
| Acquisition channels | 4 |
| Sets of data channels per packet | 2 |
| Data channels per set | 23 (46 data channels total) |
| Preamble length for balanced channel usage (msec) | 17.4 |

Table 3

[0097] Table 4, below, illustrates two example embodiments of the disclosed scaling concept that are applicable to a 25 channel FHSS system. FCC regulations require that, when 25 channels are used, the power output from the transceiver cannot exceed 250 mW. The first example in Table 4 uses four acquisition channels and 21 data channels for 10 msec data packets. In this example, the preamble length is 1.91 msec, and equal channel utilization is achieved after 84 data packets from a given transceiver. The second example, by comparison, also uses 10 msec data packets, but uses three acquisition channels and 22 data channels.

Reducing the number of acquisition channels from four to three allows the preamble length to be decreased to 1.36 msec, and equal channel utilization is achieved after 66 data packets from a given transceiver. As with the examples in Tables 1 and 2, the preamble length and number of data packets involved in achieving equal channel utilization both scale with the number of dedicated acquisition channels.

| | | |
|---|------|------|
| Data packet length (msec) | 10 | 10 |
| Total number of channels | 25 | 25 |
| Acquisition channels | 4 | 3 |
| Data channels | 21 | 22 |
| Preamble length for balanced channel usage (msec) | 1.91 | 1.36 |

Table 4

[0098] As can be appreciated from the examples of Tables 1-4, the number of acquisition channels and the number of data channels are scalable based on several configurable parameters. This scalability allows systems embodying the disclosed subject matter to be very agile.

[0099] As a particular example, the ability to scale the numbers of acquisition channels and data channels allows the system to include variable data rates and variable modulation schemes without impacting the equal channel utilization requirements of the FCC. If the network issues a preamble at one rate and data at another rate, the packet split of acquisition channels to data channels can be adjusted as needed. Similarly, the preamble modulation type does not have to adhere to the data modulation type because bit timing can be adjusted.

[0100] Another example embodiment of a system that has scalable numbers of acquisition channels and data channels includes a limited amount of data within the acquisition portion of the message. That is, a small amount of data that is associated with acquisition may be transmitted on the acquisition channels. This acquisition data may include, for example, encryption information, IP designations, particular routing information, acknowledgement requirements, and the like. Including a payload associated with the acquisition portion of the message allows more flexibility to be built into the messaging concept, while still scaling the ratios of acquisition channels to data channels to maintain equal channel utilization.

[0101] According to another aspect of the disclosed subject matter, scaling of the ratios of acquisition channels to data channels may be implemented in battery-powered devices, such as water and gas transceivers, which operate with less available energy. Reducing the

number of acquisition channels and, thus, the preamble lengths of self-coherent messages allows flexibility in adapting to such devices. Further, battery usage can be reduced, and product life can be increased, for example, to 20 years or longer.

[0102] Normally, a water or gas device only wakes up to transmit periodically. For example, such a device might wake up to transmit every four hours and listen for data updates immediately after transmitting. These battery-powered devices are normally associated with one or more electric meters that are powered full time from the AC line. To minimize system traffic, only a limited number of electric meters receive data from the battery-powered devices.

[0103] In conventional systems, long preamble times for battery-powered devices are problematic because the total packet length is increased, thereby increasing the power required to transmit the packet. Reducing the number of acquisition channels allows the total message length to be reduced, which in turn reduces the power required to transmit the packet. Battery-powered devices typically operate from two different crystals: one for sleep and another for communication when awake. The sleep crystal tolerance is typically +/- 100 ppm, which means that over a four-hour sleep interval, there is the potential for approximately +/- 4 seconds of timing error between the electric meter and the battery-powered gas or water meter.

[0104] If the gas meter and/or water meter use their own sets of acquisition channels, given the aforementioned crystal tolerance of +/- 100 ppm, the associated electric meter can scan the gas or water meter acquisition channels for the 8 seconds (*i.e.*, +/- 4 seconds) of potential battery-powered transmissions. In this way, contention within the system can be reduced while maintaining a short preamble length. After the electric meter receives the gas or water data, it can resynchronize the battery-powered device while it acknowledges receipt of the metering data. The electric meter can then spend the four-hour sleep interval, except for a few seconds, listening for electric meter traffic. The electric meter only needs to dedicate a small window of exclusive channel scans for gas and/or water meter traffic. As a result, system performance can be improved because the dedicated acquisition channels will have very little traffic.

[0105] It will be appreciated that the four acquisition channels used for a gas or water device can be in addition to the 50 channels used for the electric device. The 900 MHz ISM spectrum can be divided into as many as 65 channels that are each 400 kHz wide. FCC regulations only require that, for 1 W operation, a minimum of 50 channels must be used. For 250 mW operation, FCC regulations require the use of a minimum of 25 channels. The

disclosed embodiments facilitate the use of a greater number of channels while still maintaining a relatively short preamble length.

[0106] According to another aspect of the disclosed subject matter, a system may employ a coarse synchronization approach in which communication channels are changed in a pseudo-random sequence periodically, *e.g.*, every 400 msec. This coarse synchronization approach requires that a receiver look for both the expected channel and the possible channels on either side of the expected channel. This is because it is assumed that the system clocks may drift between time synchronization updates.

[0107] For example, an asynchronous system may use 25 communication channels at 250 mW maximum power in a coarsely synchronized network. All devices are set to the same channel for a given 400 msec window. At the conclusion of the window, all devices move to the next channel in the pseudo-random sequence.

[0108] If four dedicated acquisition channels are added to this coarsely synchronized network, a hybrid synchronous/asynchronous system can be realized in which electric devices are coarsely synchronous with the control node and battery-powered devices can be asynchronous. In this hybrid system, a total of seven channels would be scanned for a preamble. These seven channels include the expected synchronous data channel, two possible time-slipped channels on either side of the expected synchronous data channel, and the four dedicated acquisition channels. An advantage of such a hybrid system is that battery-powered devices that are not synchronized can wake up and transmit a preamble on one of the four dedicated acquisition channels, followed by data on one of the 25 communication channels. Table 5, below, discloses an example system using 25 data channels and four acquisition channels. In this system, the preamble is 8 msec to achieve equal channel utilization for a given device after transmitting 175 data packets. In this embodiment, the preamble needs to be long enough to allow receiving devices enough time to scan seven channels rather than just the four dedicated acquisition channels.

| | |
|---|-----|
| Data packet length (msec) | 50 |
| Total number of channels | 29 |
| Acquisition channels | 4 |
| Synchronous data channels | 25 |
| Asynchronous data channels | 25 |
| Preamble length for balanced channel usage (msec) | 8.0 |

Table 5

[0109] Electric powered devices that are designed for coarsely synchronous operation can transmit on the expected synchronous data channel as long as they are maintaining time through communication to a parent device in the mesh. Because the preamble and data are on the same channel for synchronous devices, the preamble time limit is not set by the equal channel utilization requirement, but rather by the time needed for a receiver to scan seven channels for the preamble. If one of the synchronous devices loses contact for a period of time, one of the four dedicated acquisition channels can be used to regain coarse time synchronization.

[0110] It will be appreciated that, in a coarsely synchronized system, time synchronization messages within the network only need to be issued every 15-30 minutes. As a result, synchronization traffic can be kept low. By comparison, ad hoc networks need to maintain sub-millisecond synchronization for system performance.

[0111] The disclosed embodiments may realize certain advantages over conventional systems. For example, by scanning a limited number of acquisition channels, rather than all frequency channels, for a carrier, less time can be spent searching for a carrier. As a result, preamble length, and overall packet length, can be reduced. Reducing overall packet length reduces the amount of communication traffic in the system, improving performance generally. Further, power consumption can be reduced, thereby particularly benefiting battery-powered devices.

[0112] These efficiencies can be realized while still achieving equal channel utilization as required by the FCC. In addition, the maximum transmission power can be increased. Some conventional systems use 25 frequency channels with a maximum power of 250 mW as permitted by the FCC. The FCC permits a maximum power of 1 W if the number of frequency channels is increased to 50. If, however, the number of frequency channels in a conventional system is increased to 50, the preamble length must also be increased because twice as many channels need to be scanned for a carrier. By comparison, limiting the number of acquisition channels that need to be scanned, as taught by the disclosed subject matter, permits a maximum power of 1 W without needing to increase the preamble length.

[0113] While systems and methods have been described and illustrated with reference to specific embodiments, those skilled in the art will recognize that modification and variations may be made without departing from the principles described above and set forth in the following claims. For example, although in the embodiments described above, the systems and methods of the present invention are described in the context of a network of metering devices, such as electricity, gas, or water meters, it is understood that the present invention can

be implemented in any kind of network in which it is necessary to obtain information from or to provide information to end devices in the system, including without limitation, networks comprising meters, in-home displays, in-home thermostats, load control devices, or any combination of such devices. Accordingly, reference should be made to the following claims as describing the scope of the present invention.

What is claimed is:

1. A wireless network comprising:
a control node; and
a plurality of communication nodes in wireless communication with the control node, each of the communication nodes having a wireless communication path to the control node that is either a direct path or an indirect path through one or more other communication nodes that serve as repeaters,
a communication node communicating a communication packet comprising a preamble portion communicated on a dedicated acquisition channel of a frequency spectrum and a data portion communicated on at least one dedicated data channel of the frequency spectrum.
2. The wireless network of claim 1, wherein the frequency spectrum is divided into a set of dedicated acquisition channels and a set of dedicated data channels.
3. The wireless network of claim 2, wherein a number of dedicated acquisition channels is selected to minimize a length of the preamble portion while maintaining substantially equal channel utilization.
4. The wireless network of claim 2, wherein at least one communication node is a battery-powered device, and wherein a subset of the set of dedicated acquisition channels is dedicated to the battery-powered device.
5. The wireless network of claim 2, wherein a first number of dedicated acquisition channels and a second number of dedicated data channels are selected as a function of a length of the communication packet to maintain substantially equal channel utilization.
6. The wireless network of claim 1, wherein the preamble portion and the data portion of the communication packet are communicated at different data rates or with different modulation schemes.
7. The wireless network of claim 1, wherein the plurality of communication nodes comprises an electrical device in coarse synchronization with the control node and a battery-powered device that is not in synchronization with the control node.

8. In a network having a control node that communicates with a plurality of communication nodes, wherein each of the communication nodes has a wireless communication path to the control node that is either a direct path or an indirect path through one or more other communication nodes that serve as repeaters, a method of operating the network, the method comprising:

scanning a plurality of dedicated acquisition channels of a frequency spectrum to detect a preamble portion of a communication packet;

ranking the dedicated acquisition channels in descending order of respective received signal strength indicator (RSSI) values;

attempting to detect a carrier on at least some of the dedicated acquisition channels; and

receiving a data portion of the communication packet using at least one dedicated data channel distinct from the dedicated acquisition channels.

9. The method of claim 8, wherein a number of dedicated acquisition channels is selected to minimize a length of the preamble portion while maintaining substantially equal channel utilization.

10. The method of claim 8, wherein at least one communication node is a battery-powered device, and wherein a subset of the set of dedicated acquisition channels is dedicated to the battery-powered device.

11. The method of claim 8, wherein a first number of dedicated acquisition channels and a second number of dedicated data channels are selected as a function of a length of the communication packet to maintain substantially equal channel utilization.

12. The method of claim 8, further comprising communicating the preamble portion and the data portion of the communication packet at different data rates or with different modulation schemes.

13. The method of claim 8, wherein the plurality of communication nodes comprises an electrical device in coarse synchronization with the control node and a battery-powered device that is not in synchronization with the control node.

14. The method of claim 8, wherein the frequency spectrum is an Industrial, Scientific, and Medical (ISM) frequency spectrum.

15. A processor-readable storage medium storing processor-executable instructions that, when executed by a processor, cause the processor to operate a wireless mesh network having a control node that communicates with a plurality of communication nodes, wherein each of the communication nodes has a wireless communication path to the control node that is either a direct path or an indirect path through one or more other communication nodes that serve as repeaters, the processor-executable instructions comprising instructions for:

scanning a plurality of dedicated acquisition channels of a frequency spectrum to detect a preamble portion of a communication packet;

ranking the dedicated acquisition channels in descending order of respective received signal strength indicator (RSSI) values;

attempting to detect a carrier on at least some of the dedicated acquisition channels; and

receiving a data portion of the communication packet using at least one dedicated data channel distinct from the dedicated acquisition channels.

16. The processor-readable storage medium of claim 15, wherein a number of dedicated acquisition channels is selected to minimize a length of the preamble portion while maintaining substantially equal channel utilization.

17. The processor-readable storage medium of claim 15, wherein at least one communication node is a battery-powered device, and wherein a subset of the set of dedicated acquisition channels is dedicated to the battery-powered device.

18. The processor-readable storage medium of claim 15, wherein a first number of dedicated acquisition channels and a second number of dedicated data channels are selected as a function of a length of the communication packet to maintain substantially equal channel utilization.

19. The processor-readable storage medium of claim 15, the processor-executable instructions comprising instructions for communicating the preamble portion and the data portion of the communication packet at different data rates or with different modulation schemes.

20. The processor-readable storage medium of claim 15, wherein the plurality of communication nodes comprises an electrical device in coarse synchronization with the control node and a battery-powered device that is not in synchronization with the control node.

FIG. 1


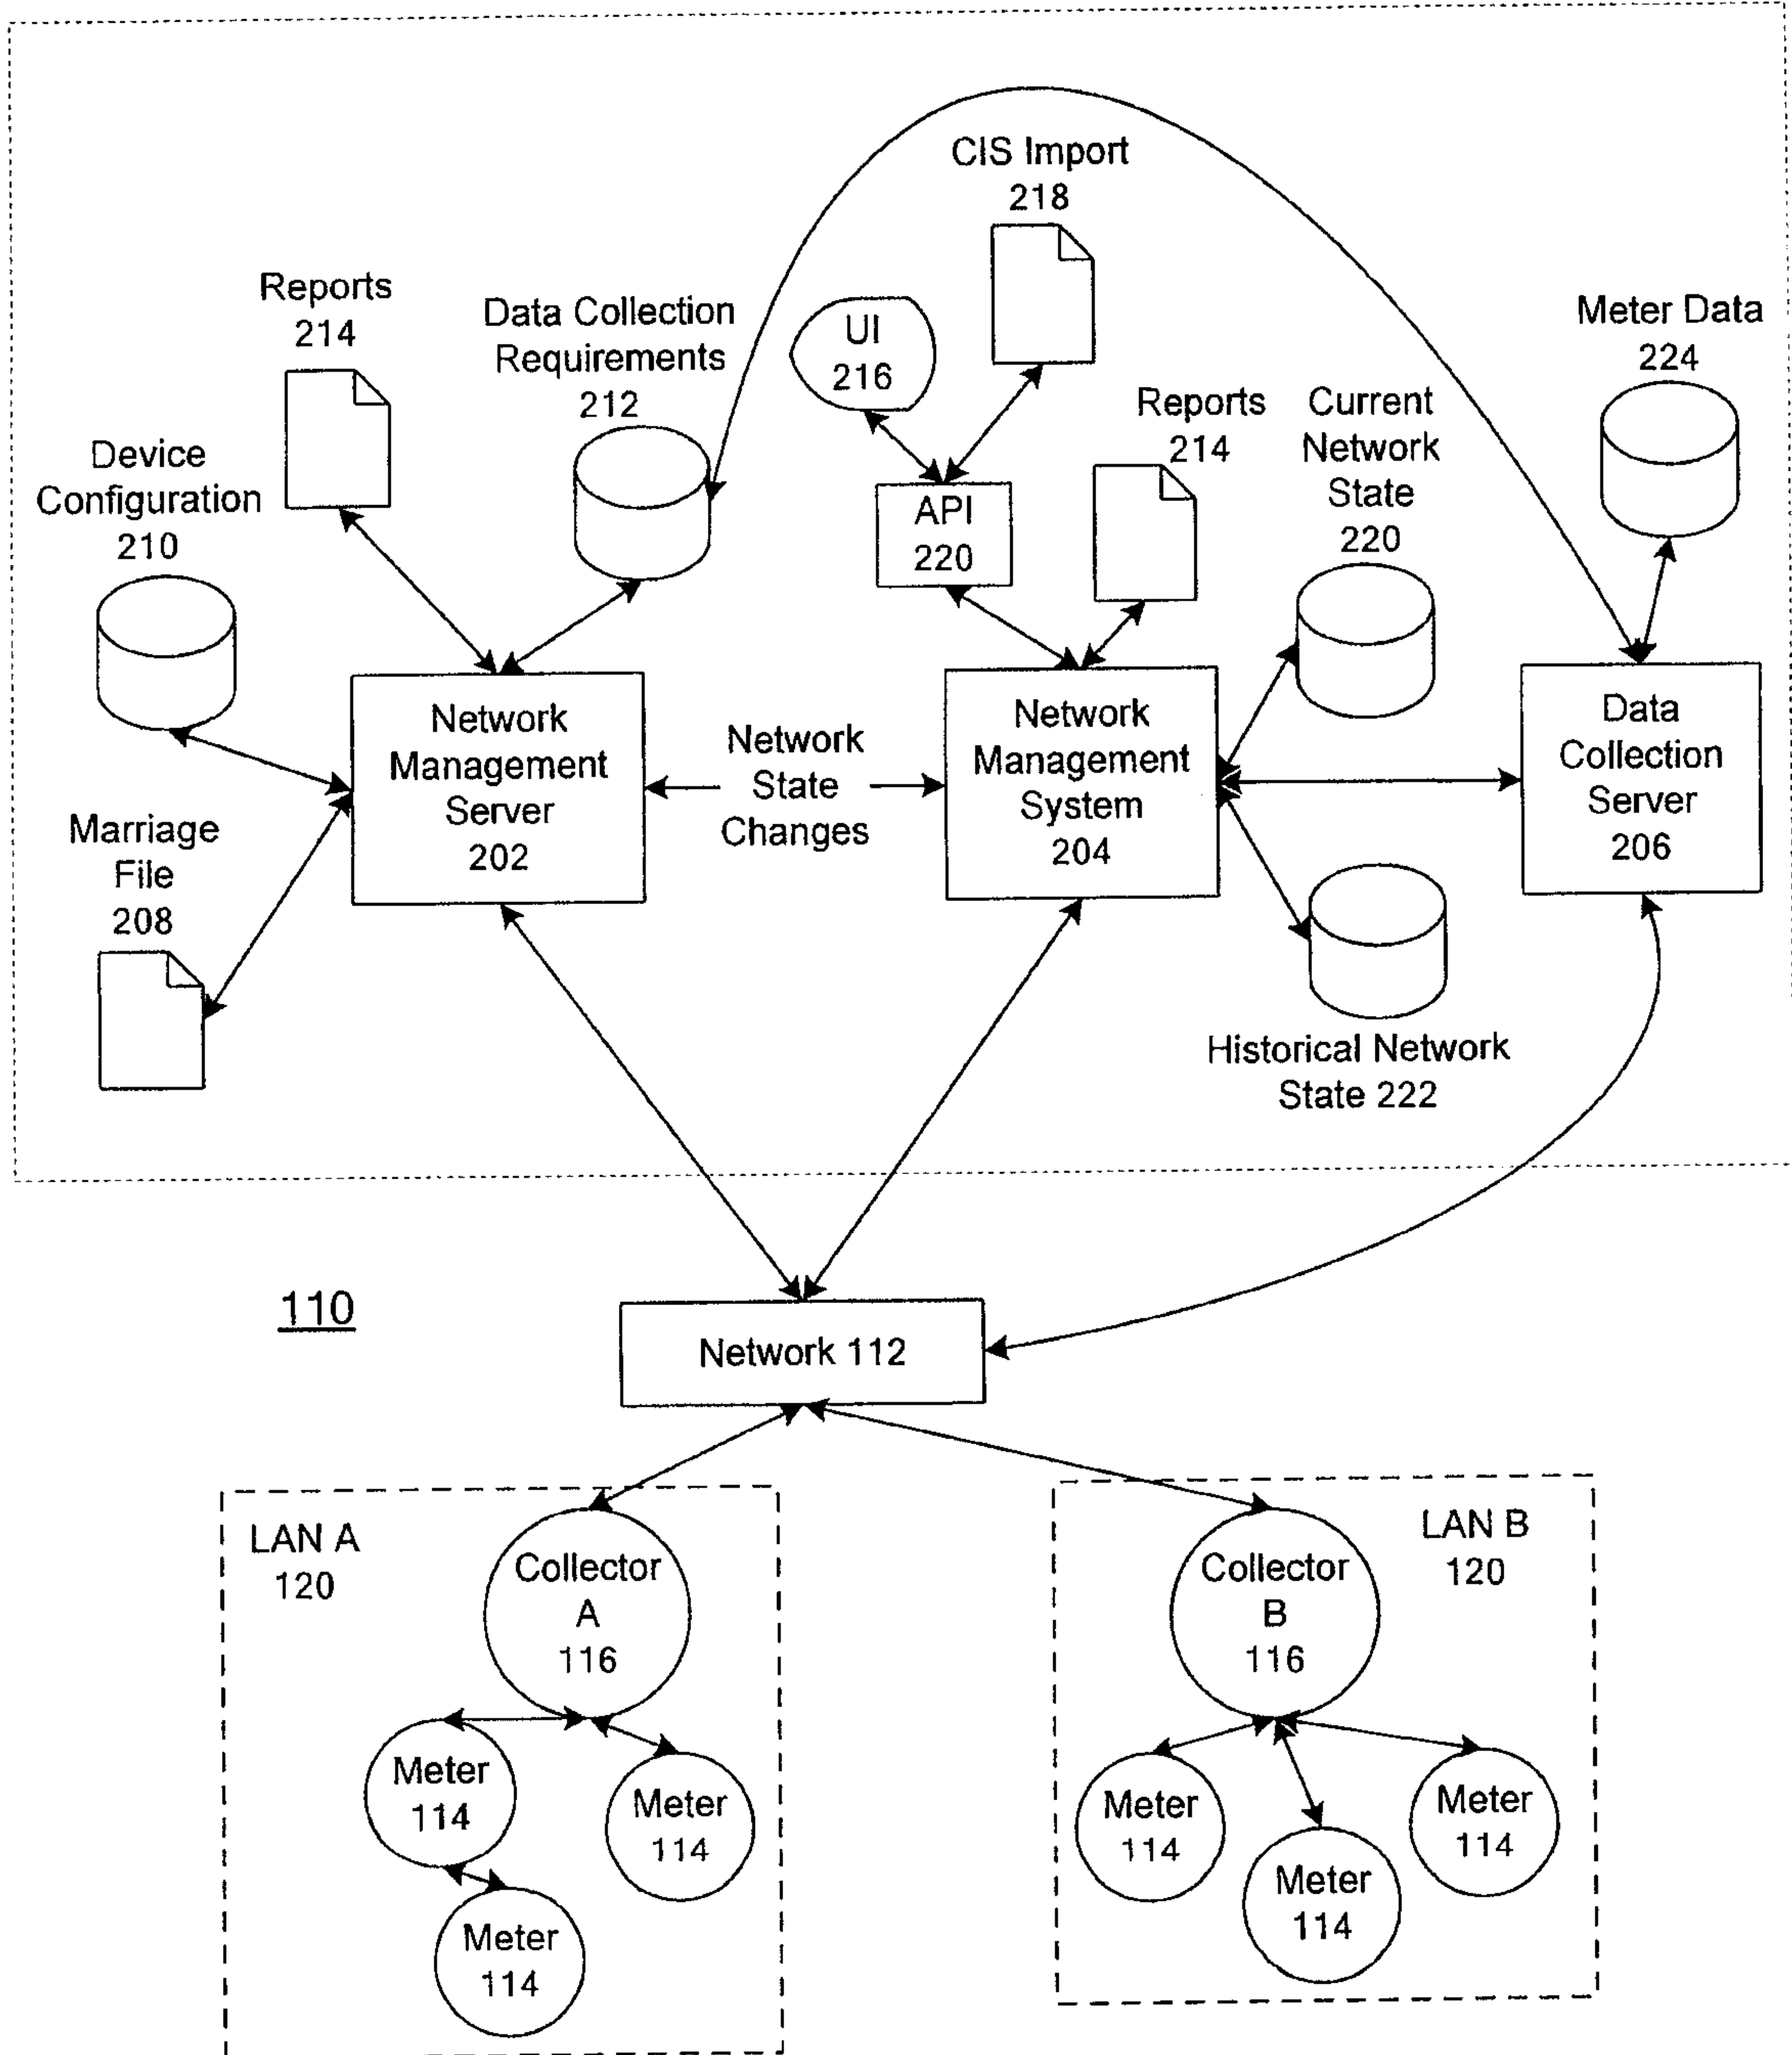


FIG. 2






Figure 3A




Figure 3B





FIG. 5

120

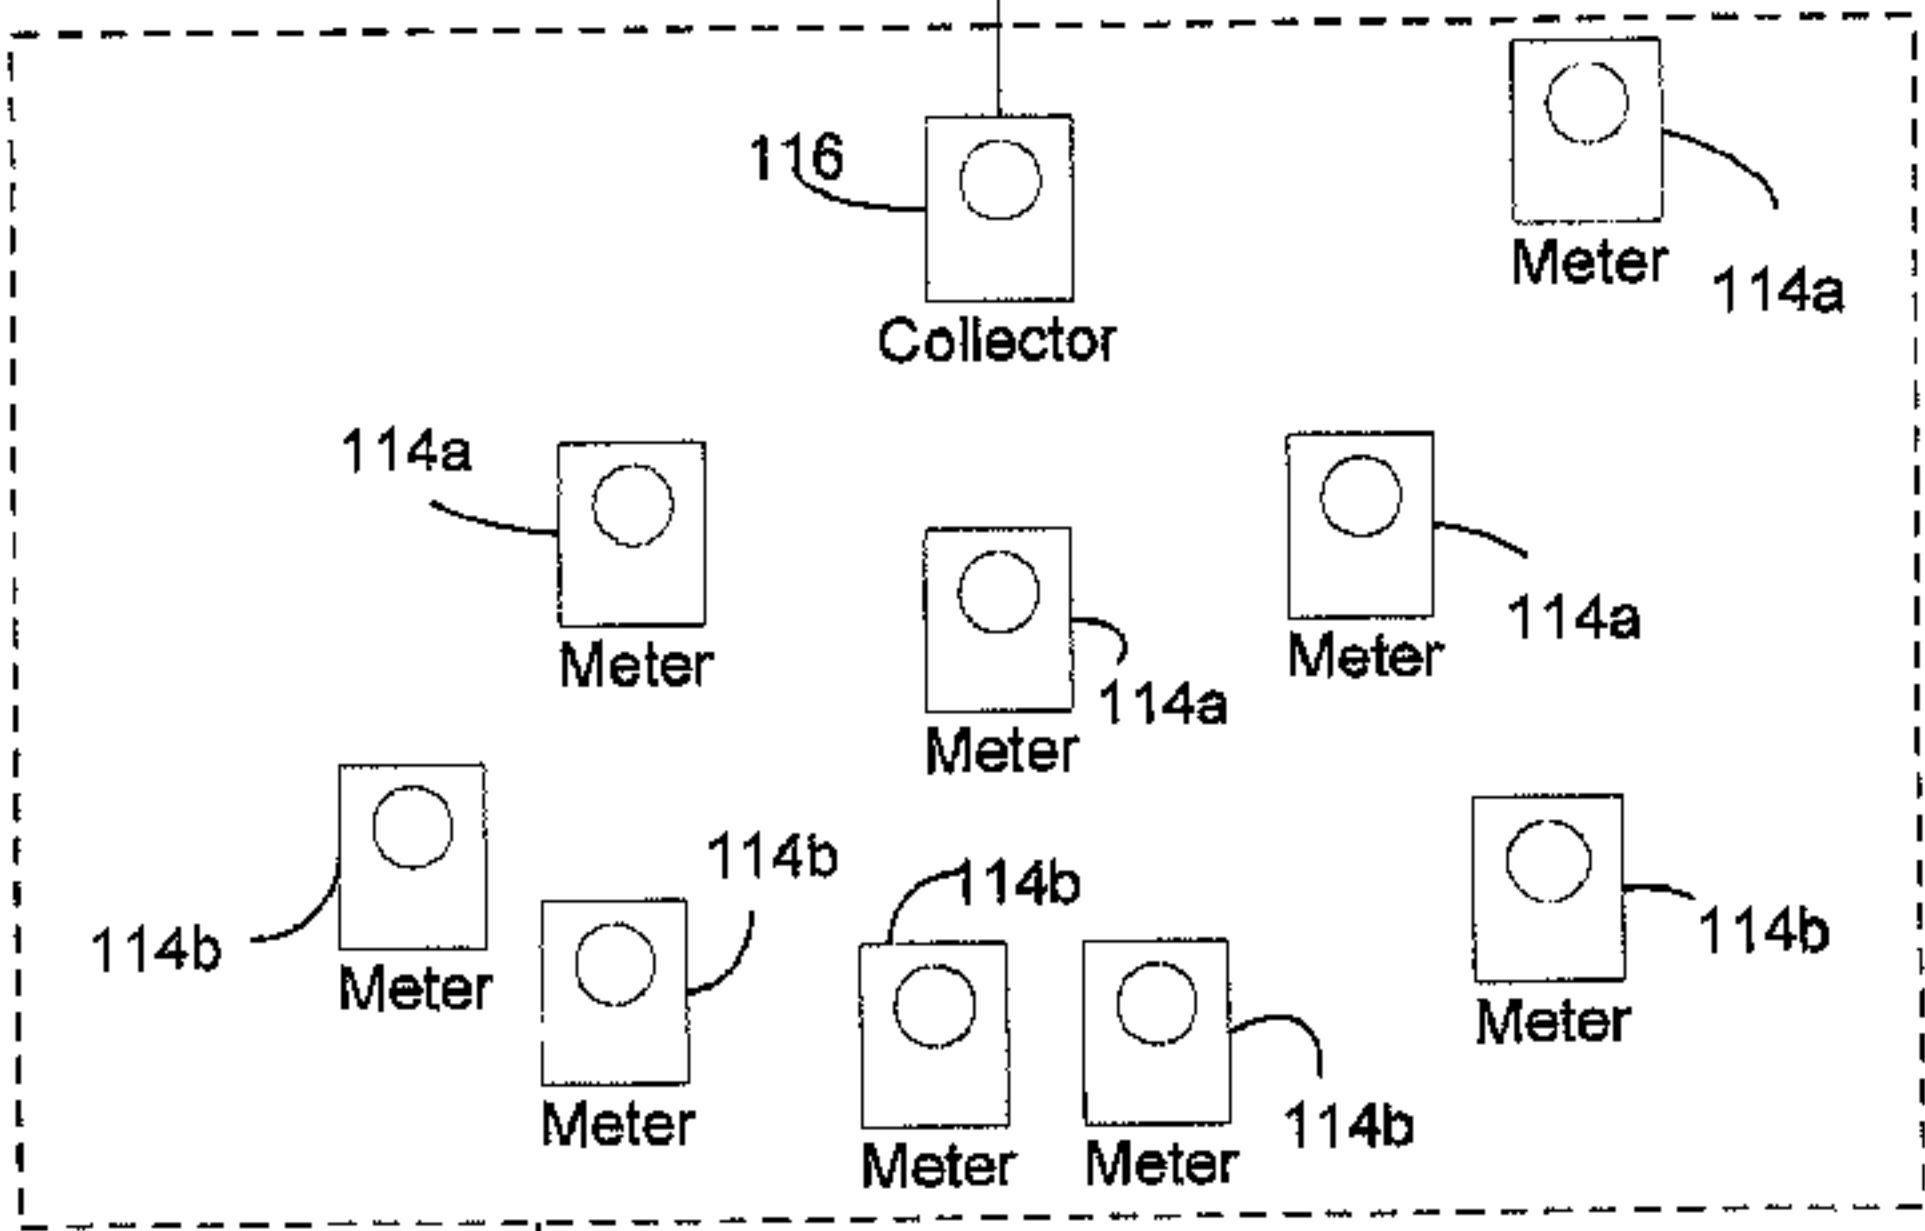


112

Data Collection Server 206



110



120