



US 20110188444A1

(19) **United States**(12) **Patent Application Publication**
Borleske et al.(10) **Pub. No.: US 2011/0188444 A1**(43) **Pub. Date: Aug. 4, 2011**(54) **HIGH PRIORITY DATA READS FOR
ACQUISITION OF REAL-TIME DATA IN
WIRELESS MESH NETWORK****Publication Classification**(51) **Int. Cl.**
H04B 7/14 (2006.01)(52) **U.S. Cl.** **370/315**(57) **ABSTRACT**

A wireless mesh network, method, and processor-readable storage medium for operating a network that has a control node that communicates with a plurality of communication nodes are disclosed herein. The control node schedules time-limited sample windows and only collects data that is designated as having a priority, such as voltage data, during those scheduled sample windows. Because the sample window has a defined end point time, the reading software that executes at the head end and that reads data from the collector can follow a known schedule because, by definition, the collector will already be finished collecting priority data when the sample window expires.

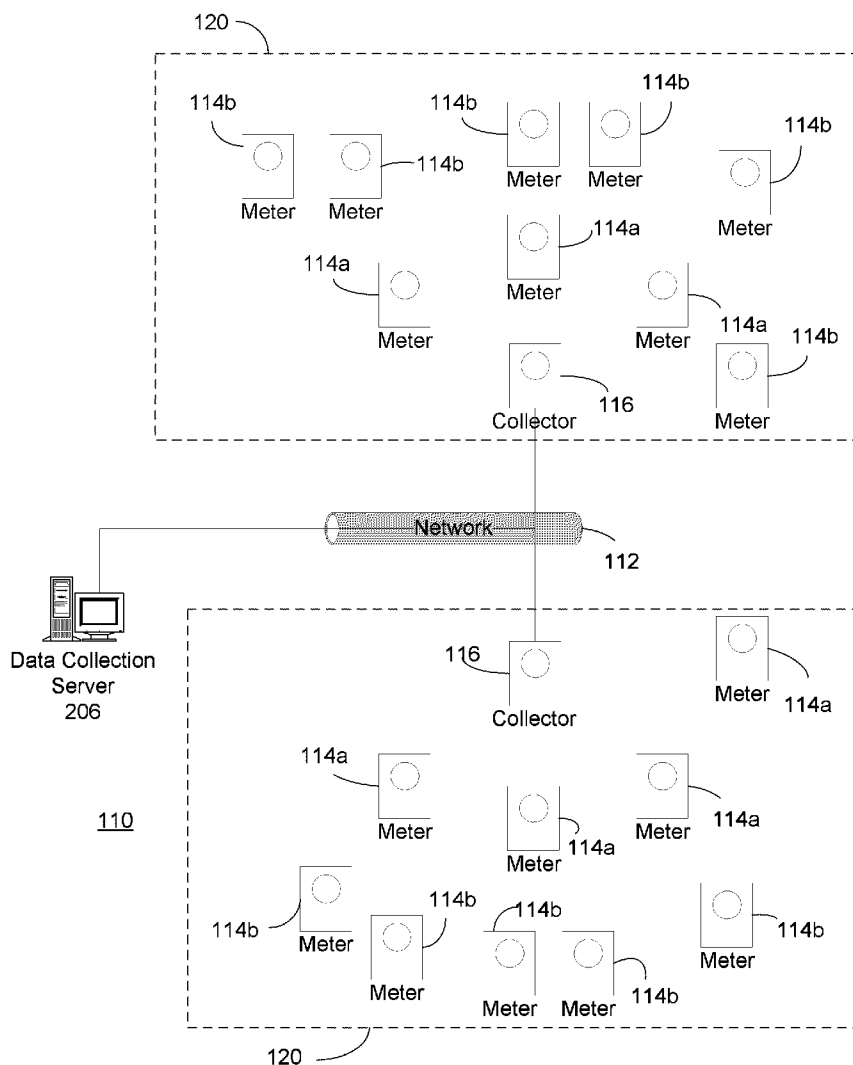
(75) **Inventors:** **Andrew J. Borleske**, Garner, NC
(US); **Robert T. Mason, JR.**,
Raleigh, NC (US)(73) **Assignee:** **ELSTER SOLUTIONS, LLC**,
Raleigh, NC (US)(21) **Appl. No.:** **13/015,780**(22) **Filed:** **Jan. 28, 2011****Related U.S. Application Data**(60) Provisional application No. 61/299,583, filed on Jan.
29, 2010.

FIG. 1

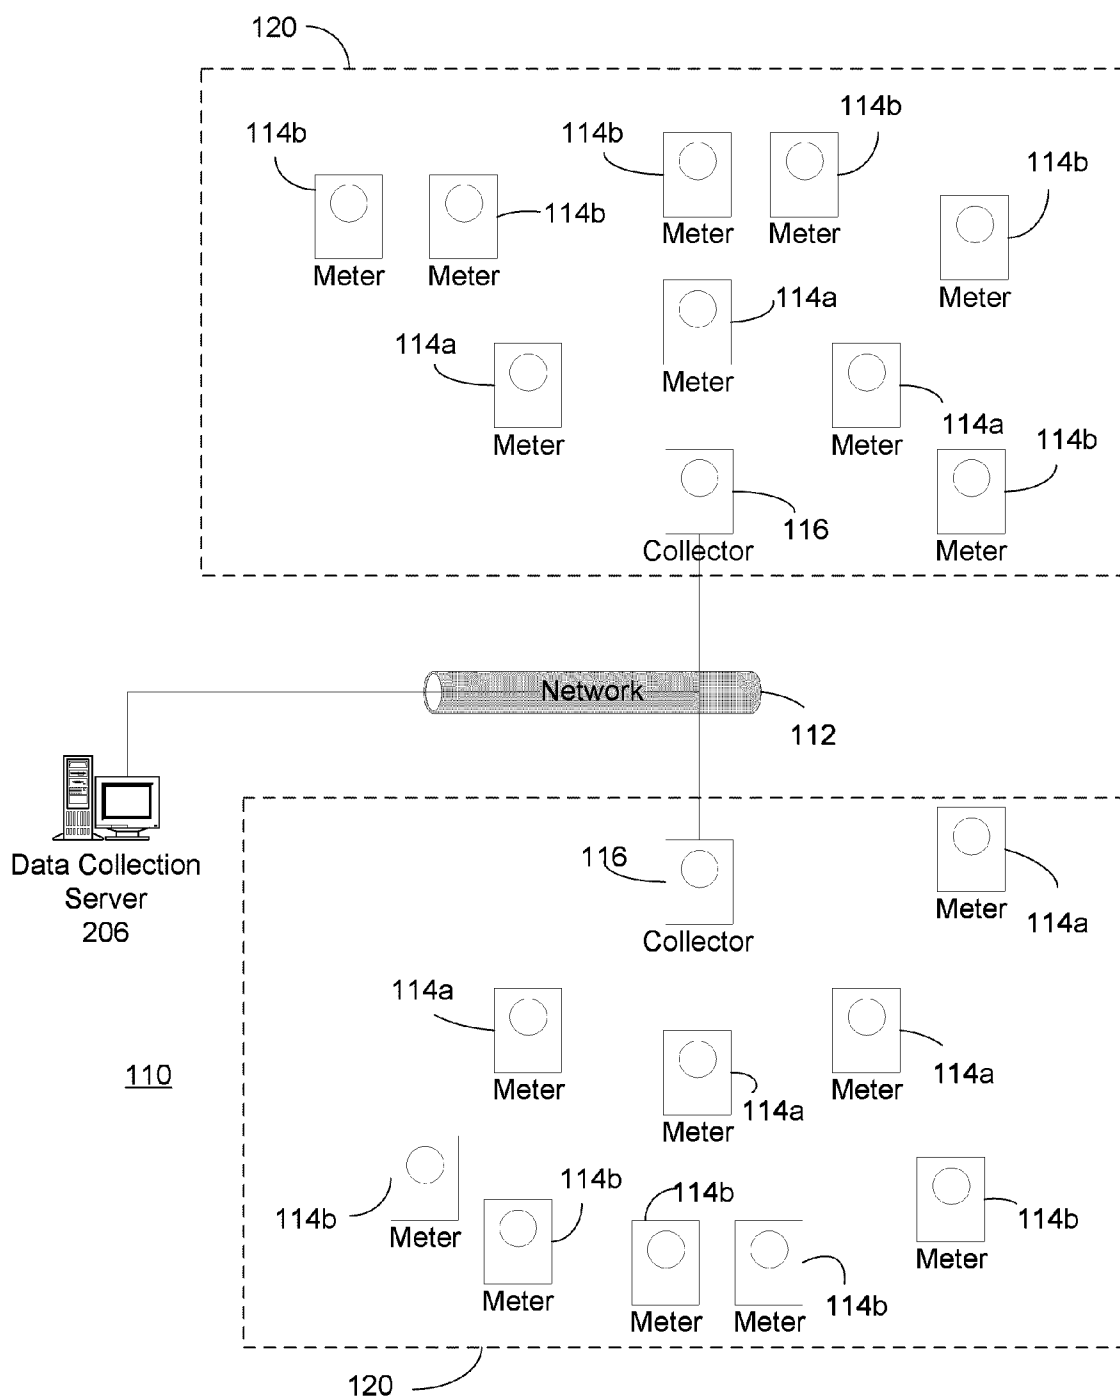
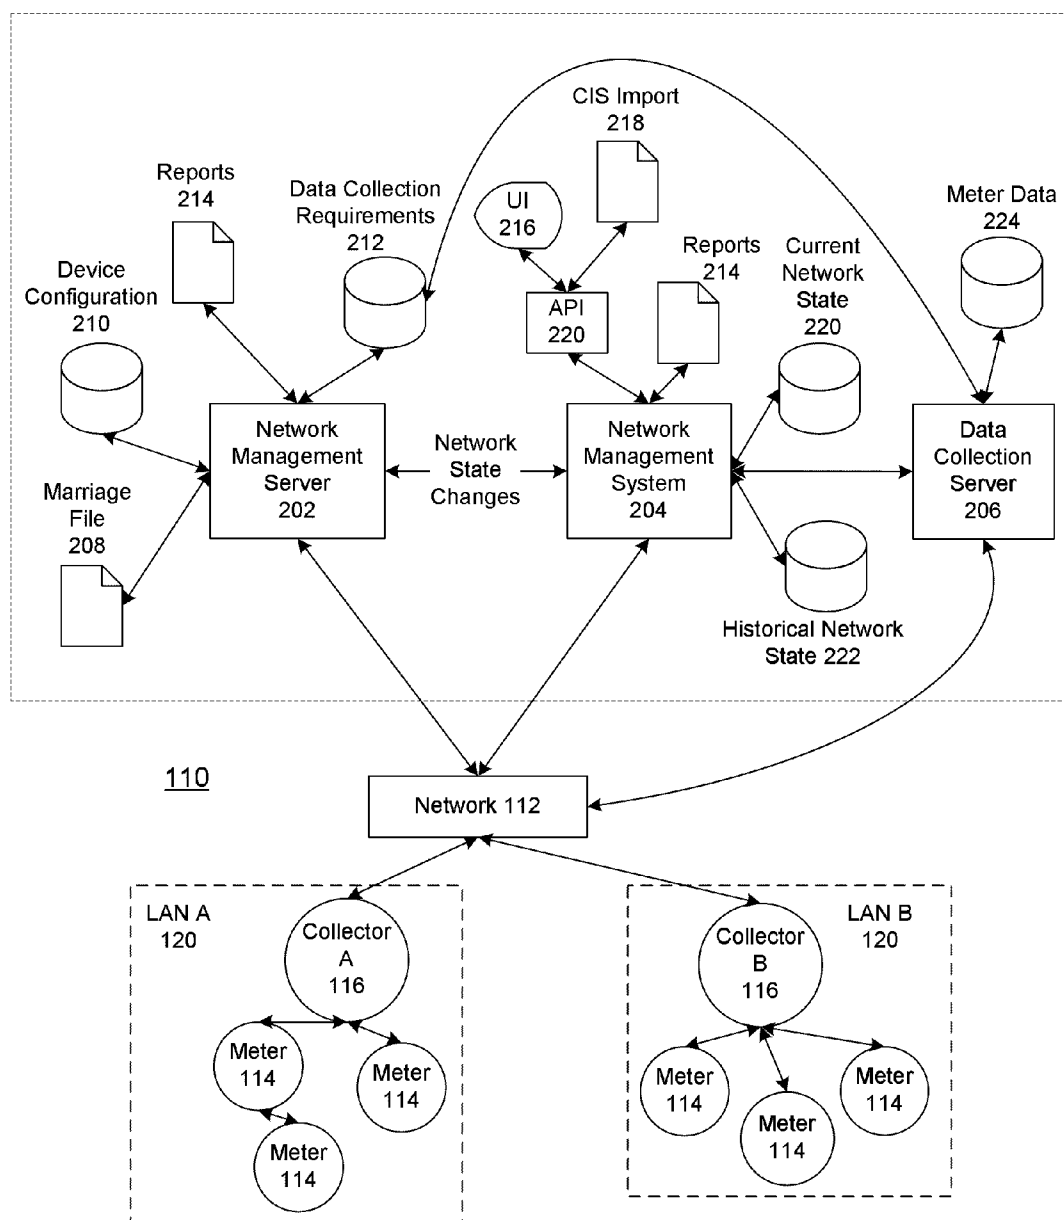


FIG. 2



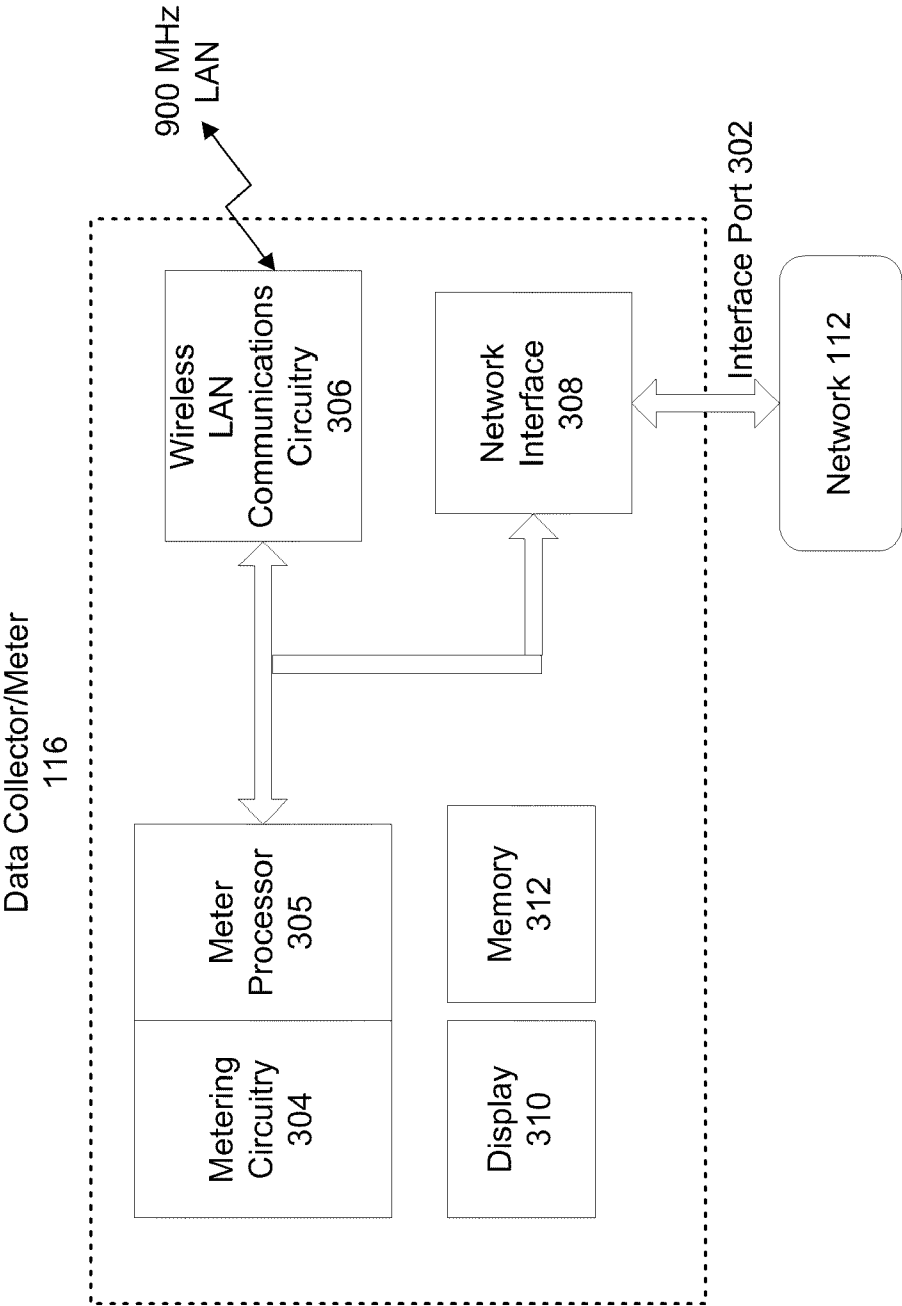


Figure 3A

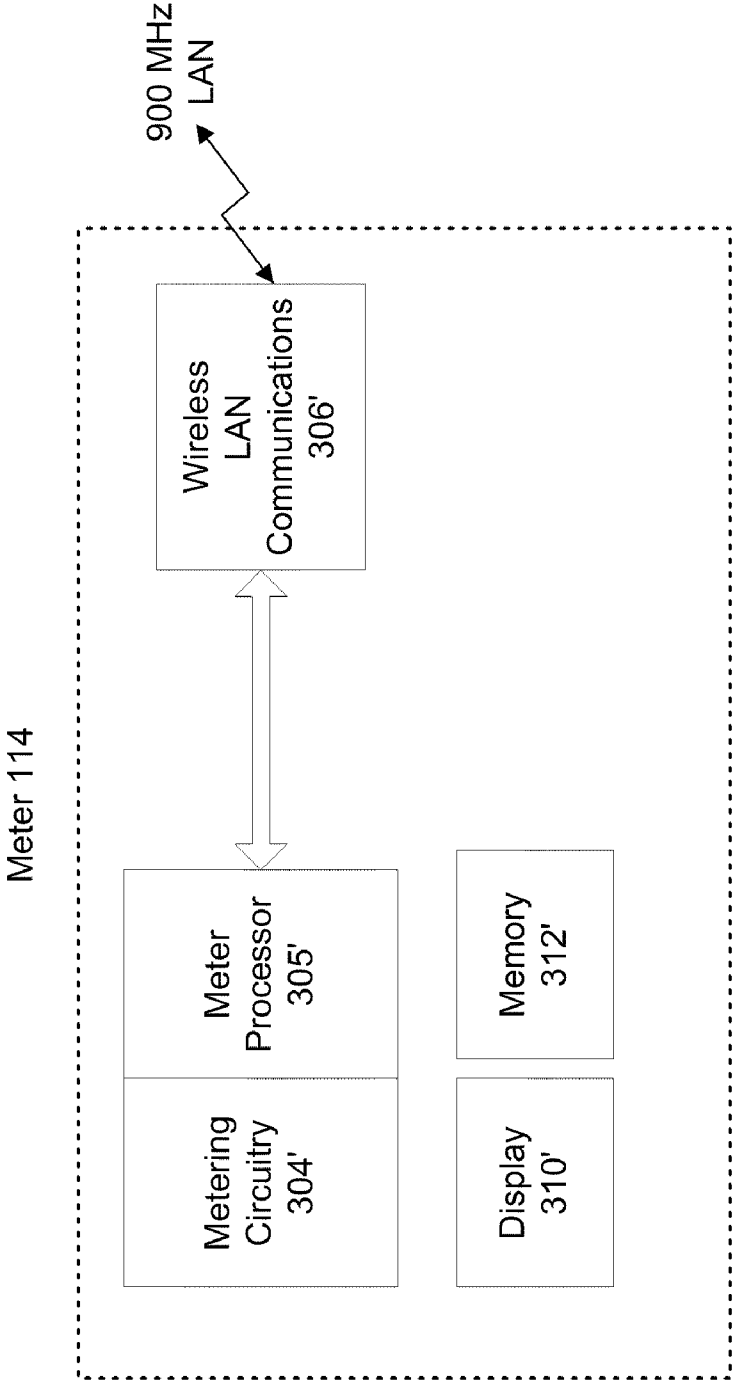
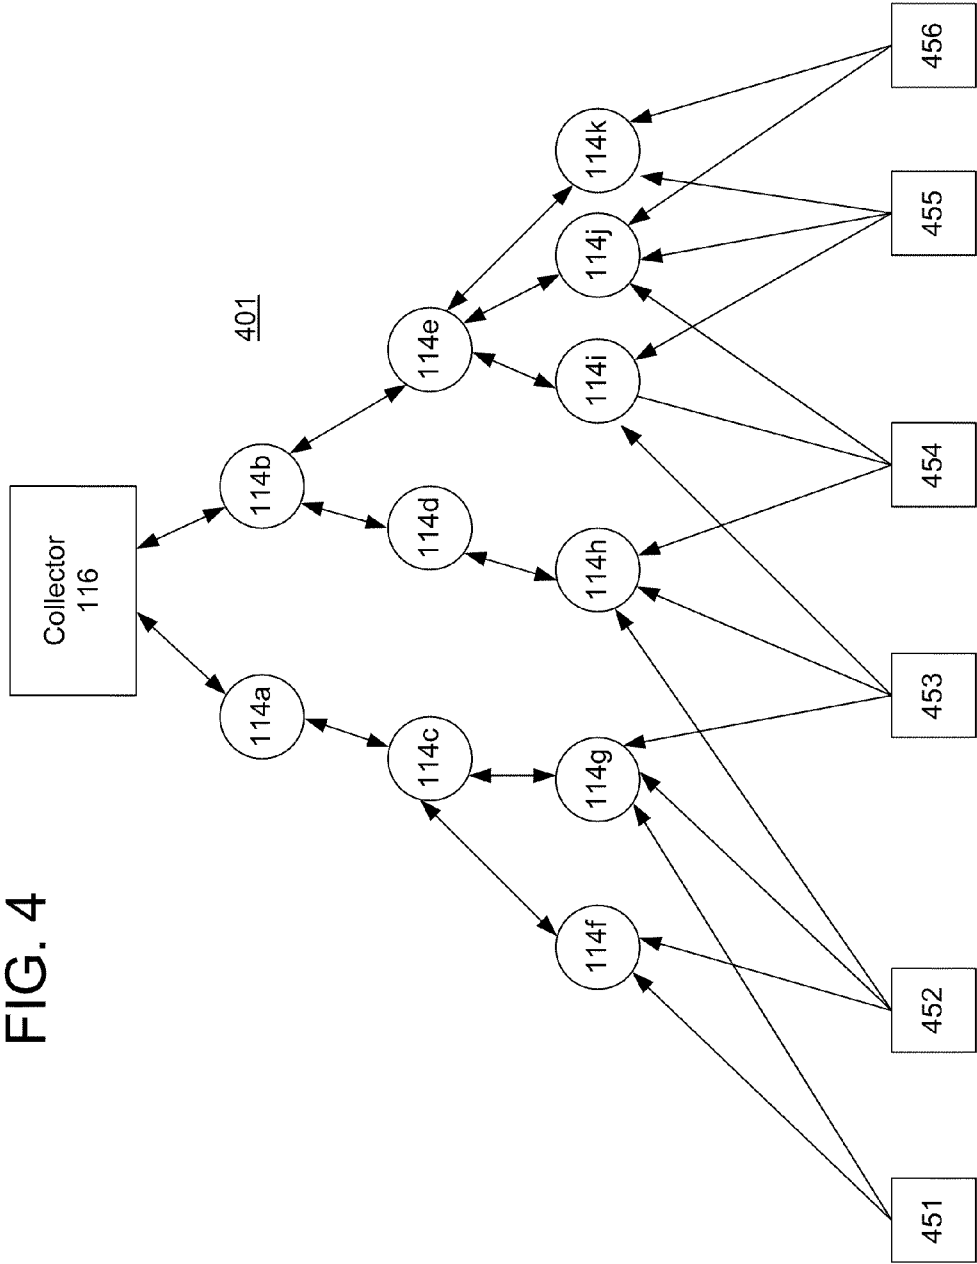


Figure 3B



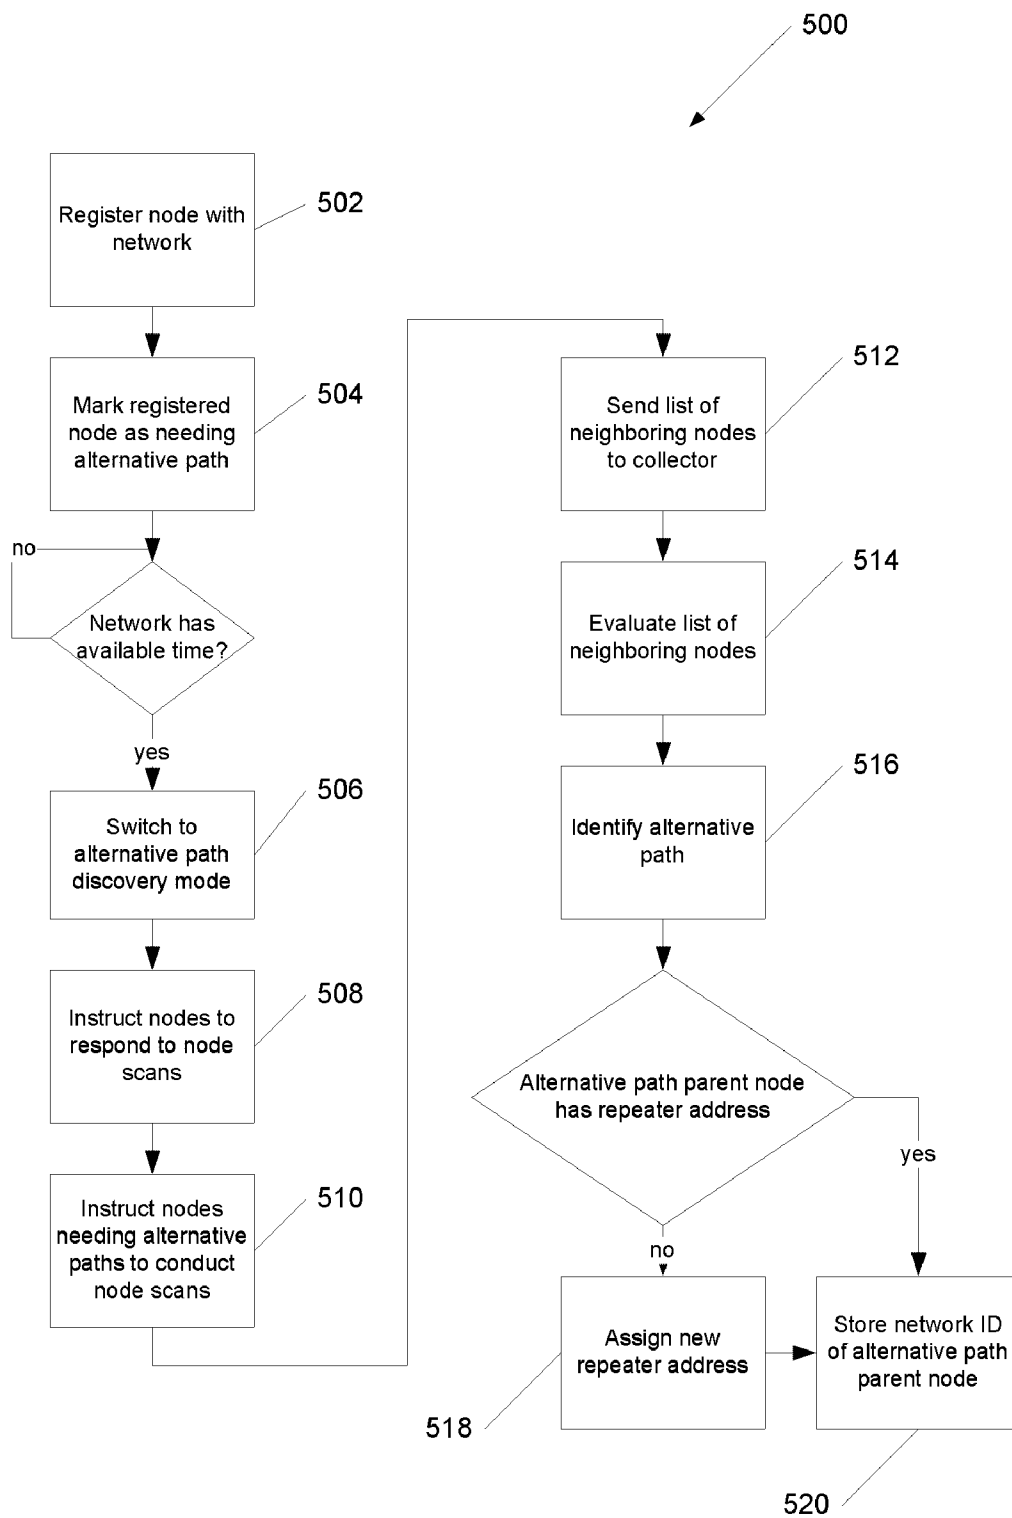


FIG. 5

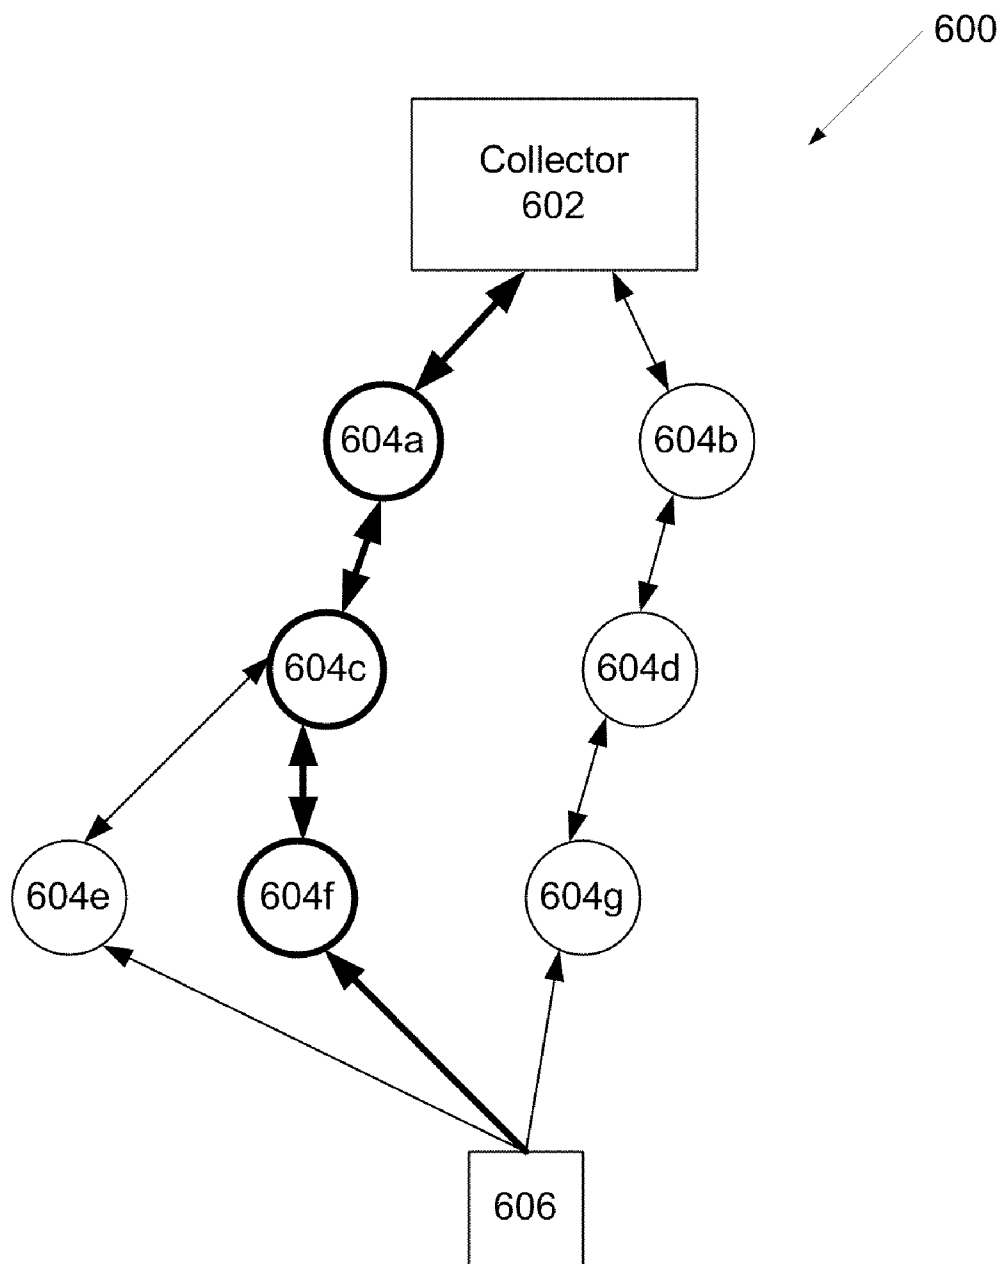


FIG. 6

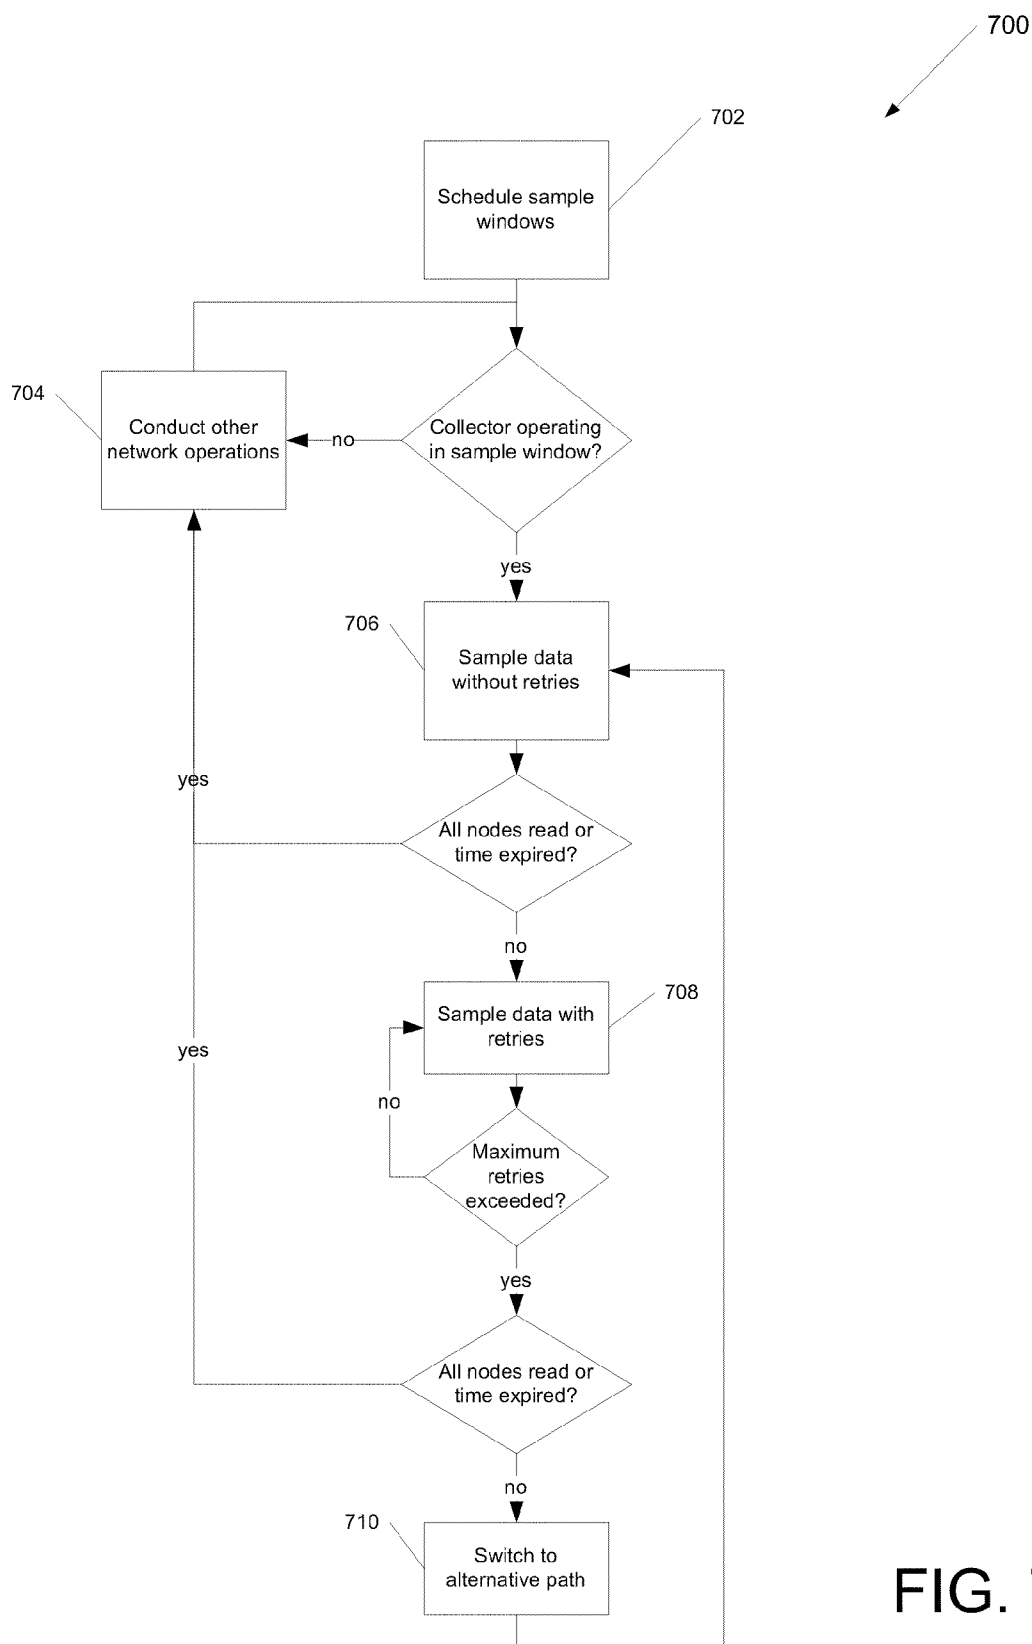


FIG. 7

HIGH PRIORITY DATA READS FOR ACQUISITION OF REAL-TIME DATA IN WIRELESS MESH NETWORK

CROSS-REFERENCE TO RELATED APPLICATION

[0001] This application claims the benefit of priority to Provisional Application Ser. No. 61/299,583, filed Jan. 29, 2010.

TECHNICAL BACKGROUND

[0002] 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.

[0003] 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. Radio technology in the 902-928 MHz frequency 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.

[0004] Some conventional utility meter reading communication networks use multiple repeaters. In some such networks, a head end computer system keeps track of the distribution network configuration and configures the repeaters to work within that distribution network. Such networks work well as long as the power line network do not change configuration and the propagation characteristics of the network remain constant. Alternate routing schemes can be built into the hardware so that the master computer can select different paths to achieve end device communication. One drawback of this type of network is that the network must be managed from the master computer. As the number of end points grows, there is a need to move intelligence further into the network and have routing decisions made in the remote hardware.

[0005] 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 which serve as repeaters for the meter data of the transmitting node.

[0006] 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.

[0007] Within these smart mesh networks, communications are achieved from a central collector through repeaters to endpoints, and the number of repeaters in a chain can be quite large. There are two different methods for extracting data from mesh networks: polling and bubble up. In a polling approach, the route from a collector to an endpoint is established and data is pulled from the endpoint by sending a unicast packet to the endpoint and back. In a bubble up approach, the data may be originated at the endpoint based on a schedule or a prior instruction, and the route to the collector or gateway can be dynamically determined.

[0008] In a controlled wireless mesh network, there are certain parameters that a network user may want to monitor on a sample or subset of nodes on a very rapid sample period. A good example of this would be voltage data available from electricity meters used by an electric utility. While the utility needs to know voltage at defined points in the system, it typically does not need it from all points. If the voltage data can be captured at a rapid rate, it can be used as a feedback mechanism for the distribution network. However, if the voltage data is collected rapidly through normal communication channels, the data collection process could potentially limit other activities on the network. Accordingly, a need exists for a way to collect data with rapid sample periods, while avoiding significantly adverse effects on other network activities.

SUMMARY OF THE DISCLOSURE

[0009] A wireless mesh network, method, and processor-readable storage medium for operating a network that has a control node that communicates with a plurality of communication nodes are disclosed herein. According to various embodiments, the control node schedules time-limited sample windows and only collects data that is designated as having a priority, such as voltage data, during those scheduled sample windows.

[0010] 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. The control node is configured to schedule sample time windows for collecting data that is designated as having a priority and to initiate a request to the communication nodes for the data that is designated as having a priority only during the scheduled sample time windows.

[0011] Another embodiment is directed to a method of operating a network that has a control node that communi-

cates with a plurality of communication nodes. 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 causing the control node to schedule sample time windows for collecting data that is designated as having a priority. During a scheduled sample time window, the control node attempts to collect the data from a communication node using a primary communication path without using any retry attempts. Also during the scheduled sample time window, if the communication node has not responded to the attempt to collect the data, then control node attempts to collect the data from the communication node using the primary communication path a number of times that does not exceed a maximum number of retry attempts using the primary communication path. The method may be implemented using a processor readable storage medium that stores instructions that, when executed by the processor, cause the processor to carry out the method.

[0012] Various embodiments may realize certain advantages. For example, because the sample window has a defined end point time, the reading software that executes at the head end and that reads data from the collector knows what time it is safe to read the data. The reading software can follow a known schedule because, by definition, the collector will already be finished collecting priority data, like voltage data, when the sample window expires.

[0013] Other features and advantages of the described embodiments may become apparent from the following detailed description and accompanying drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

[0014] 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:

[0015] FIG. 1 is a diagram of an exemplary metering system;

[0016] FIG. 2 expands upon the diagram of FIG. 1 and illustrates an exemplary metering system in greater detail;

[0017] FIG. 3A is a block diagram illustrating an exemplary collector;

[0018] FIG. 3B is a block diagram illustrating an exemplary meter;

[0019] FIG. 4 is a diagram of an exemplary subnet of a wireless network for collecting data from remote devices;

[0020] FIG. 5 is a flow diagram illustrating an example method for identifying an alternative communication path in a wireless mesh network;

[0021] FIG. 6 is a diagram illustrating an example wireless mesh network in which an alternative communication path is identified using the method of FIG. 5; and

[0022] FIG. 7 is a flow diagram illustrating an example method for sampling data in a network, according to one embodiment.

DETAILED DESCRIPTION OF ILLUSTRATIVE EMBODIMENTS

[0023] Exemplary systems and methods for gathering meter data are described below with reference to FIGS. 1-7. 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.

[0024] 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.

[0025] FIG. 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 Electric, LLC and marketed under the tradename REX.

[0026] 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).

[0027] 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 com-

prise 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.

[0028] Referring now to FIG. 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.

[0029] 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.

[0030] 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.

[0031] 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.

[0032] 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.

[0033] 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.

[0034] FIG. 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 FIG. 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 FIG. 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.

[0035] As shown in FIG. 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.

[0036] 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.

[0037] 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.

[0038] 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.

[0039] FIG. 3B is a block diagram of an exemplary embodiment of a meter 114 that may operate in the system 110 of FIGS. 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.

[0040] Referring again to FIG. 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.

[0041] 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.

[0042] While only three levels of meters are shown (collector 116, first level 114a, second level 114b) in FIG. 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.

[0043] 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.

[0044] 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.

[0045] 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.

[0046] 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.

[0047] 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.

[0048] 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.

[0049] 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.

[0050] 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.

[0051] 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.

[0052] 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.

[0053] 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 imme-

diated exception, collector **116** may wait a set period of time to allow for receipt of a message indicating the power outage has been corrected.

[0054] 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.

[0055] 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.

[0056] 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.”

[0057] 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.

[0058] 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 acknowledgements 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.

[0059] 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).

[0060] 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 serve 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**.

[0061] 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**.

[0062] 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.

[0063] 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.

[0064] 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.

[0065] 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.

[0066] 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.

[0067] 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**.

[0068] 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.

[0069] 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.

[0070] 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.

[0071] 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.

[0072] 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**.

[0073] 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.

[0074] 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;

[0075] 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

[0076] 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.

[0077] "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;

[0078] RptAddr—the ID of the parent node that serves as the next repeater for the sending node;

[0079] 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.

[0080] 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.

[0081] 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.

[0082] 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.

[0083] 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.

[0084] 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.

[0085] 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.

[0086] 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.

[0087] 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 unregister command to propagate to all nodes in the subnet. Ultimately, the meters register with the replacement collector using the procedure described above.

[0088] 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.

[0089] 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.

[0090] 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. FIG. 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.

[0091] 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.

[0092] As discussed above, in some wireless mesh networks, there is always one fixed path for a communication node at any given time, even though a communication node can switch its path if a current path becomes unreliable. In such networks, if there is a break in the current path, e.g., if one of the communication nodes forming the path malfunctions, the communication node will be unresponsive until the network can reconfigure to establish a new path from the communication node to the collector. For some network operations, this delay in the healing of the path can be unacceptable.

[0093] To reduce delays in healing the wireless mesh network, some embodiments disclosed herein involve a wireless

mesh network in which a collector maintains both a primary path and an alternative path to each communication node. If the primary path becomes unavailable, the collector can immediately attempt the communication on the alternative path without incurring the time penalty of reconfiguring the network.

[0094] FIG. 5 illustrates an example method 500 for operating a wireless mesh network according to one embodiment. Once a communication node is registered into the network at a step 502, for example, after being qualified by determining the reliability of communications between the communication node and the collector or central node, that communication node is marked as needing an alternative path at a step 504. When the network has available time, for example, when no higher priority communication task is scheduled, the network switches to an alternative path discovery mode at a step 506. In this mode, the collector or central node instructs all communication nodes to respond to node scans at a step 508.

[0095] In addition, at a step 510, the collector or central node instructs all communication nodes that have been marked as needing an alternative path to conduct a node scan and to qualify communication nodes that are identified in the node scan. To qualify an identified communication node, the communication node that conducts the node scan evaluates the quality of the link between it and the identified communication node. The identified communication node is qualified if the quality of the link meets a performance criterion. For example, the communication node that conducts the node scan may repeatedly send a message that requires a response to the identified communication node and wait for a response from the identified communication node. Signal strength may affect the link quality, but other factors may also enter this determination. For instance, if interfering devices are present and communicate using the ISM band, the link quality may be adversely affected even though the signal strength is relatively high.

[0096] At a step 512, after a communication node has evaluated the link quality with identified neighboring communication nodes and has qualified appropriate communication nodes, it sends a list of qualified neighboring communication nodes to the collector. The collector then evaluates the list of neighboring communication nodes at a step 514.

[0097] At a step 516, the collector identifies an alternative path to the communication node that is most different from the primary path, but that still meets minimum performance qualification criteria. In particular, considering only paths from the collector to the communication node that satisfy the minimum performance qualification criteria, the collector selects, as the alternative path, the path that involves the lowest-level common parent node as the primary path. The lowest-level common parent node is defined as the communication node that is located the least number of communication hops away from the collector.

[0098] By way of example and not limitation, FIG. 6 illustrates a portion of a wireless mesh network 600 having a collector 602, communication nodes 604a-604g, and an endpoint node 606. Assuming that all of the depicted communication nodes 604a-604g are appropriately qualified, there are three communication paths from the collector 602 to the endpoint node 606. One path, the primary path indicated by bold circles and bold arrows in FIG. 6, comprises communication nodes 604a, 604c, and 604f. In addition, there is a second path comprising communication nodes 604a, 604c,

and 604e. There is also a third path comprising communication nodes 604b, 604d, and 604g.

[0099] In the example wireless mesh network 600, the second path shares a common parent node with the primary path, namely, communication node 604c. That is, communication node 604c is common to both the primary path and the second path. As shown in FIG. 6, communication node 604c is located two hops away from the collector 602.

[0100] The third path shares only the collector 602 itself as a common parent node with the primary path. Because the collector 602 is a higher level common parent node than communication node 604c, the third path is deemed the most different from the primary path and is selected as the alternative path.

[0101] The alternative path is formed by the existing primary path to the new alternative path parent node, e.g., communication node 604g in the example shown in FIG. 6, plus the last hop from the alternative path parent device to the endpoint device. Accordingly, in the example of FIG. 6, the alternative path comprises communication nodes 604b, 604d, and 604g. Referring back to FIG. 5, if a neighboring communication node is designated as a parent on an alternative path and is not already assigned a network repeater address, the collector assigns a new repeater address to the alternative path parent node at a step 518. Thus, the alternative path parent node can be used as a router to the endpoint device. Because the alternative path parent node is also registered to the same collector, the collector can store the network identifier associated with the alternative path parent node as the alternative path address for the endpoint device at a step 520. The collector uses the primary path to the alternative path parent node and the last hop from that node to the endpoint device as the full alternative communication path.

[0102] For any high priority network task, if communication on the primary path fails, the collector may automatically switch to the alternative path. Since the alternative path parent node has already been discovered and made a repeater at step 518 of FIG. 5, the switch between the primary path and the alternative path consumes a negligible amount of time.

[0103] In some embodiments, if communications on the primary path repeatedly fail, the collector may permanently designate the alternative communication path as a new primary communication path to a communication node and may repeat the process of FIG. 5 to find a new alternative communication path for the end device. Failure of communications may be defined, for example, as a configurable number of successive communication failures or as a communication success rate that falls below a threshold.

[0104] As discussed above, there are certain parameters that a network user would want to monitor on a sample or subset of nodes on a very rapid sample period, such as voltage data available from electricity meters used by an electric utility. While the utility needs to know voltage data at defined points in the system, it typically does not need the data from all points. If voltage data can be captured at a rapid rate, it can be used as a feedback mechanism for the distribution network. However, if a normal communication channel is used to obtain this data, the data collection process may adversely affect other activities on the network.

[0105] To facilitate obtaining data at a rapid rate without adversely affecting other network activities, the network may use a defined communication window for certain communications that are designated as having priority over other com-

munications. For example, messages communicating voltage readings may be designated as high priority.

[0106] FIG. 7 illustrates an example method 700 for sampling data according to one embodiment. At a step 702, the collector schedules defined, time limited sample windows for sampling data. The sample windows are scheduled such that only a fixed percentage of network time can be devoted to them. For example, the amount of time that can be devoted to the sample windows may be limited to 50% of the total network time.

[0107] During operation of the network, when the collector is not operating in a scheduled sample window, it may conduct other network operations, as shown at a step 704. If the collector is operating in a scheduled sample window, on the other hand, the collector initiates a request for priority data from the communication nodes and begins sampling data from the communication nodes without using application level retries at a step 706 to collect as much data as possible, as quickly as possible. Once the collector has progressed through the list of communication nodes, the collector then attempts to sample data from any communication nodes that did not respond to the previous sampling attempt at a step 708. The collector can retry sampling data from communication nodes up to a number of attempts that is configurable, e.g., four retry attempts.

[0108] When the collector has exhausted all of the retry attempts, the collector may employ alternative communication paths to attempt to collect data from any nodes that have not yet responded to the collector at a step 710. These alternative communication paths may be discovered, for example, using the method 500 shown in FIG. 5. It will be appreciated by those skilled in the art that, while FIG. 7 depicts retry attempts being conducted first on the primary communication path at step 708 before any attempts are performed using the alternative communication path at step 710, these steps can be performed in other sequences. For instance, the collector may switch between the primary communication path and the alternative communication path with each retry attempt, essentially performing step 710 before step 708.

[0109] The process of steps 706-710 continues until either all of the nodes have been read, or until the time limit of the sample window expires. Because the sample window has a defined end point time, the reading software that executes at the head end and that reads data from the collector knows what time it is safe to read the data. The reading software can follow a known schedule because, by definition, the collector will already be finished collecting priority data, like voltage data, when the sample window expires. In some embodiments, to further reduce latency, the collector can notify the head end directly when the data is available. In this way, if all of the data is read before the sample window expires, the head end can retrieve the data from the collector even though the sample window has not yet expired. Because both the end point collection times and head end data collection times are known, the data can be published directly to the consumer and treated as "fresh" data to be used as feedback on the system.

[0110] 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 elec-

tricity, 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,

wherein the control node is configured to schedule sample time windows for collecting data that is designated as having a priority and to initiate a request to the communication nodes for the data that is designated as having a priority only during the scheduled sample time windows.

2. The wireless network of claim 1, wherein the control node is configured to, during a scheduled sample time window:

attempt to collect the data from a communication node using a primary communication path without using any application level retry attempts; and

if the communication node has not responded to the attempt to collect the data, then attempt to collect the data from the communication node using the primary communication path a number of times that does not exceed a maximum number of application level retry attempts using the primary communication path.

3. The wireless network of claim 2, wherein the control node is further configured to, during the scheduled sample time window:

if the communication node has not responded to the attempt to collect the data using the primary communication path, then attempt to collect the data from the communication node using an alternative communication path a number of times that does not exceed a maximum number of attempts using the alternative communication path.

4. The wireless network of claim 1, wherein the control node is configured to, during a scheduled sample time window:

attempt to collect the data from a communication node using a primary communication path;

if the communication node has not responded to the attempt to collect the data using the primary communication path, then attempt to collect the data from the communication node using an alternative communication path; and

if the communication node has not responded to the attempts to collect the data using the primary communication path and the alternative communication path, then attempt to collect the data from the communication node, alternating between using the primary communication path and using the secondary communication path a number of times that does not exceed a maximum number of attempts.

5. The wireless network of claim 1, wherein a head end device reads the collected data from the control node when the scheduled sample time window expires.

6. The wireless network of claim 5, wherein the control node notifies the head end device if the control node has collected data from all of the communication nodes in wireless communication with the control node.

7. The wireless network of claim 1, wherein the data that is designated as having a priority comprises voltage data.

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 causing the control node to:

schedule sample time windows for collecting data that is designated as having a priority;

during a scheduled sample time window, attempt to collect the data from a communication node using a primary communication path without using any application level retry attempts; and

during the scheduled sample time window, if the communication node has not responded to the attempt to collect the data, then attempt to collect the data from the communication node using the primary communication path a number of times that does not exceed a maximum number of application level retry attempts using the primary communication path.

9. The method of claim 8, further comprising causing the control node to, during the scheduled sample time window:

if the communication node has not responded to the attempt to collect the data using the primary communication path, then attempt to collect the data from the communication node using an alternative communication path a number of times that does not exceed a maximum number of attempts using the alternative communication path.

10. The method of claim 8, further comprising causing the control node to, during a scheduled sample time window:

attempt to collect the data from a communication node using a primary communication path;

if the communication node has not responded to the attempt to collect the data using the primary communication path, then attempt to collect the data from the communication node using an alternative communication path; and

if the communication node has not responded to the attempts to collect the data using the primary communication path and the alternative communication path, then attempt to collect the data from the communication node, alternating between using the primary communication path and using the secondary communication path a number of times that does not exceed a maximum number of attempts.

11. The method of claim 8, further comprising using a head end device to read the collected data from the control node when the scheduled sample time window expires.

12. The method of claim 11, further comprising causing the control node to notify the head end device if the control node has collected data from all of the communication nodes in wireless communication with the control node.

13. The method of claim 8, wherein the data that is designated as having a priority comprises voltage data.

14. The method of claim 8, further comprising repeating, for other communication nodes of the plurality of communication nodes, the steps of causing the control node to attempt to collect the data from a communication node using a primary communication path without using any retry attempts and causing the control node to attempt to collect the data from the communication node using the primary communication path a number of times that does not exceed a maximum number of retry attempts using the primary communication path.

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 causing the control node to:

schedule sample time windows for collecting data that is designated as having a priority;

during a scheduled sample time window, attempt to collect the data from a communication node using a primary communication path without using any application level retry attempts; and

during the scheduled sample time window, if the communication node has not responded to the attempt to collect the data, then attempt to collect the data from the communication node using the primary communication path a number of times that does not exceed a maximum number of application level retry attempts using the primary communication path.

16. The processor-readable storage medium of claim 15, further storing processor-executable instructions for causing the control node to, during the scheduled sample time window:

if the communication node has not responded to the attempt to collect the data using the primary communication path, then attempt to collect the data from the communication node using an alternative communication path a number of times that does not exceed a maximum number of attempts using the alternative communication path.

17. The processor-readable storage medium of claim 15, further storing processor-executable instructions for causing the control node to, during a scheduled sample time window:

attempt to collect the data from a communication node using a primary communication path;

if the communication node has not responded to the attempt to collect the data using the primary communication path, then attempt to collect the data from the communication node using an alternative communication path; and

if the communication node has not responded to the attempts to collect the data using the primary communication path and the alternative communication path, then attempt to collect the data from the communication node, alternating between using the primary communication path and using the secondary communication path a number of times that does not exceed a maximum number of attempts.

18. The processor-readable storage medium of claim **15**, further storing processor-executable instructions for causing a head end device to read the collected data from the control node when the scheduled sample time window expires.

19. The processor-readable storage medium of claim **18**, further storing processor-executable instructions for causing the control node to notify the head end device if the control

node has collected data from all of the communication nodes in wireless communication with the control node.

20. The processor-readable storage medium of claim **15**, wherein the data that is designated as having a priority comprises voltage data.

* * * * *