



- (51) International Patent Classification:
H04L 12/26 (2006.01) H04L 12/24 (2006.01)
- (21) International Application Number:
PCT/US2014/016538
- (22) International Filing Date:
14 February 2014 (14.02.2014)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
61/766,551 19 February 2013 (19.02.2013) US
61/779,222 13 March 2013 (13.03.2013) US
13/896,029 16 May 2013 (16.05.2013) US
- (71) Applicant: **POWER TAGGING TECHNOLOGIES, INC.** [US/US]; 5425 Airport Boulevard, Boulder, CO 80301 (US).
- (72) Inventors: **BERNHEIM, Henrik, F.**; 7777 East 1st Place #109, Denver, CO 80230 (US). **HANSELL, Jerritt**; 4625 Pitkin Drive, Boulder, CO 80303 (US). **MARTIN, Marcia, Reid**; 1209 S. Emery St., Longmont, CO 80501 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.


(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Published:

— with international search report (Art. 21(3))

[Continued on next page]

(54) Title: METHODS FOR ANALYZING AND OPTIMIZING THE PERFORMANCE OF A DATA COLLECTION NETWORK OF AN ELECTRICAL DISTRIBUTION GRID



(57) Abstract: A system and methods for optimizing the performance of communication network utilizing an electrical distribution grid are disclosed. Optimization methods may include archiving historical message data and metrics from transmissions by Remote Hubs received on one or more Substation-to-Edge channels. Trends in the archived metrics over time are later analyzed to determine which combination of transmission variables such as frequency band, modulation method, drive voltage, and transmission time produce the highest message success rates overall. The results of such analysis may be used to provide feedback to the Remote Hubs or reveal necessary repairs to the network. Optimization may also be carried out locally by an individual Remote Hub by estimating the impedance of the transmission medium and adjusting either drive voltage or transmission band.

WO 2014/130365 A1

- *before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments (Rule 48.2(h))*

METHODS FOR ANALYZING AND OPTIMIZING THE PERFORMANCE OF A DATA COLLECTION NETWORK ON AN ELECTRICAL DISTRIBUTION GRID

RELATED APPLICATIONS

[0001] This application claims priority to U.S. Provisional Application No. 61/766,551, filed on February 19, 2013, and U.S. Provisional Application No. 61/779,222, filed on March 13, 2013, the contents of which are incorporated herein by reference.

FIELD OF THE INVENTION

[0002] The present invention is directed toward employing the electrical distribution grid as a short and long-range transmission medium and data-bearing network, and further toward the use of signals and messages on the network for the purpose of inferring schematic and topological properties of the distribution grid, which vary over time.

BACKGROUND OF THE INVENTION

[0003] The power grid is generally considered to be composed of two logical regions, the Transmission Grid(s) and the Distribution Grid(s). The Transmission Grid originates at large generation points such as hydroelectric dams, nuclear reactors, wind farms, and coal-fired or gas-fired power plants. Power from the generation point is transmitted as high-voltage alternating current (AC) over a loosely connected network of long, high-voltage lines to points where demand for power exists, such as factories, farms, and population centers. At the edges of the Transmission Grid there is a collection of Distribution Substations. Distribution Substations contain one or more Substation Transformers, which step down the voltage from high transmission line levels (typically 130kV to 700kV) to the medium voltage levels (typically from 4kV to about 35kV) at which power is distributed to consumers within a distribution service area. At the edge of the Distribution Grid are a number of Service Transformers, which transform the medium voltage of the distribution grid to the low voltages (in the US, typically 120V, 208V, 240V, 277V, or 480V). Other voltages in addition to some of these can be used elsewhere in the world. In some cases, a tier of one or more transformers, called step-down transformers, lying schematically between the Substation Transformers and the Service Transformers, create intermediate voltage reductions between the Substation and

the Service Transformers. Each Service Transformer powers one or more metered loads. A load can be a dwelling, a commercial or industrial building, an element of municipal infrastructure such as a series of street lamps, or agricultural apparatus such as irrigation systems. A typical distribution grid includes other elements used to balance and regulate the flow of power. Examples of such elements are capacitor banks, voltage regulators, switches, and reclosers. Figure 1 illustrates a typical segment of the power grid.

[0004] Distribution grids have been designed and deployed in a variety of topological configurations. In the United States, distribution grid types are typically characterized as radial, loop, or networked. Other emerging cases are the campus grid and the microgrid. Additional topologies, not described, are used elsewhere in the world.

[0005] Figure 2a is a topological schematic of a typical radial grid. In a radial grid, a substation has one or more substation transformers. Each substation transformer has one or more substation busses. One or more three-phase feeders “radiate” outward from each substation bus, with single-phase, double-phase, or three-phase lateral lines branching off from the feeders, and tap-off points (or simply “taps”) in turn branching from the laterals. Radial grids are inexpensive to design and build because they are simple, but they are most vulnerable to outages because they lack redundant power paths, so that any break causes at least one load to lose power.

[0006] Figure 2b is a topological schematic of a typical loop distribution grid. In a loop grid, each end of select feeders is attached to a power source such as a bus of a substation transformer. If the loop is undamaged, then power is available at all loads if either substation transformer is operational. If there is a break in the loop, then power is available at all loads assuming that both transformers are operational. In normal circumstances, a system of switches is used to ensure that only one substation transformer at a time is delivering power to each segment of the grid.

[0007] Figure 2c is a topological schematic of a typical networked grid. This topology has maximum redundancy. In addition to employing multiple power sources, all the service transformers are linked to one another on the secondary side in a mesh arrangement. Multiple breaks in connectivity are required to cause a power outage at any point. Networked grids are most expensive to build and maintain, and are typically used

in major urban areas such as Manhattan or Washington, DC where high-value, high-criticality loads are concentrated together.

[0008] Figure 2d shows a microgrid or campus network. Microgrids are not traditional in electrical distribution technology, but are emerging as a response to increased focus on energy conservation and on distributed generation of energy from renewable sources. Many variations are possible. This type of grid is typically attached to, but severable from, a wider distribution grid, and may contain its own power sources such as windmills, solar panels, or rechargeable storage batteries as well as loads. The entire network may employ low-voltage lines.

[0009] A distribution substation receives high-voltage power from the transmission grid into one or more large power transformers. A distribution transformer may incorporate a type of regulator called a load-tap changer, which alters the voltage the transformer delivers to a power distribution bus (the substation bus) by including or excluding some turns of the secondary winding circuit of the transformer, thereby changing the ratio of input to output voltage. One or more feeders depend from the substation bus. If too many feeders are required, additional transformers and busses are used.

[0010] In order to monitor and control the components of the grid, current transformers (CTs) or other current sensors such as Hall-effect sensors are attached to power-bearing conductors within the substation. The CTs output a low current on a looped conductor, which is accurately proportional to the current delivered at the high voltage conductor being monitored. These low-current outputs are suitable for connecting to data acquisition subsystems associated with Supervisory Control and Data Acquisition (SCADA) systems in the substation. Primary monitoring CTs are designed and built into the substation, because changing or adding CTs to the high-voltage components is impossible or dangerous while current is flowing. On the other hand, additional CTs may be safely added to the low-current SCADA loops as needed without impacting power delivery.

[0011] In addition to the power lines themselves, the distribution grid contains numerous other devices intended to regulate, isolate, stabilize, and divert the flow of power. These devices include switches, reclosers, capacitor banks (usually for power factor correction), and secondary voltage regulators. All these devices affect the behavior of the distribution

grid when considered as a data-bearing network, as do the various loads and secondary power sources on the grid. Devices that have abrupt state changes will introduce impulse noise on the grid, as can loads turning on and off. Some devices, such as transformers and capacitor banks, filter and attenuate signals at certain frequencies.

[0012] Other than the wires connecting a consumer load and the associated meter to a service transformer, the service transformer is the outermost element of the distribution grid before the power is actually delivered to a consumer. The meter is attached at the point where the power from the service transformer is delivered to the consumer. Service transformers can be three-phase, dual-phase, or single phase, as can meters.

[0013] Traditionally, reading meters was one of the largest operational costs incurred by electrical utilities. Original electric meters were analog devices with an optical read-out that had to be manually examined monthly to drive the utility billing process. Beginning in the 1970s, mechanisms for digitizing meter data and automating its collection began to be deployed. These mechanisms evolved from walk-by or drive-by systems where the meter would broadcast its current reading using a short-range radio signal, which was received by a device carried by the meter reader. These early systems were known as Automated Meter Reading systems or AMRs. Later, a variety of purpose-built data collection networks, employing a combination of short-range RF repeaters in a mesh configuration with collection points equipped with broadband backhaul means for transporting aggregated readings began to be deployed.

[0014] These networks were capable of two-way communication between the “metering head-end” at a utility service center and the meters at the edge of this data collection network, which is generally called the Advanced Metering Infrastructure or AMI. AMIs can collect and store readings frequently, typically as often as every 15 minutes, and can report them nearly that often. They can read any meter on demand provided that this feature is used sparingly, and can connect or disconnect any meter on demand as well. AMI meters can pass signals to consumer devices for the purpose of energy conservation, demand management, and variable-rate billing. Because the AMI network is separate from the power distribution grid, AMI meters are neither aware of nor sensitive to changes in the grid topology or certain conditions on the grid. Nonetheless, the introduction of AMI is generally considered to be the beginning of the Smart Grid.

[0015] Many characteristics of the electrical distribution infrastructure have limited the success of efforts to use the grid itself as a communications medium. First, the grid is a noisy environment. As already noted, state changes in loads on the grid as well as control and regulation artifacts on the grid itself cause impulse noise on the power line. Normal operation of loads like electrical motors, simple variations in the overall load, and ambient RF noise (chiefly from lightening and other weather-related causes) add up to significant Gaussian noise. The measured noise floor at a typical substation in the United States sits at about 80-90dB below the maximum amplitude of the 60 Hz fundamental. The complex impedance of the grid varies across both the frequency and time domains. This may lead to loss of signal at a receiver sited at a higher voltage point on the grid when impedance increases, or alternately force the transmitter to use more energy than would be necessary on the average. Capacitor banks sited at points along the grid for the purpose of optimizing the power factor can cause signal attenuation. Most significantly, transformers act as low-pass filters, dramatically attenuating signals above a certain frequency. The threshold frequency is not the same on every distribution grid, because different arrangements and types of transformers are employed and because the transformers themselves are not deliberately tuned to filter at a specified frequency. All these variables impact the frequency response of the medium.

[0016] Additionally, injecting modulated current signals on the grid may cause interference between the injected signals themselves. One problematic phenomenon is crosstalk, where a signal injected on one power line is detectable on another line. When crosstalk occurs on two or more phases of the same feeder, it can be caused by inductive and capacitive coupling, as the phase lines run alongside one another for most of the length of the feeder. Crosstalk may also be due to multiple phase windings on the same transformer core. Feeder-to-feeder crosstalk has also been measured, and may be caused by reflection of the injected signal off the power bus at the substation. Given the complexity, diversity, and age of the distribution grids in the United States and the world, less is known about these phenomena than might be expected.

[0017] Finally, using the distribution grid as a communications medium often has side effects which interfere with the primary purpose of the grid, which of course is delivering clean, reliable power to consumers. If devices under power resonate with an injected

current, a phenomenon called *flicker* results. LED, CFL, incandescent and fluorescent lighting visibly flickers in response to certain frequencies. This is annoying and sometimes dangerous, as visual flicker has been demonstrated to cause both seizures and vertigo. Other types of devices, such as fans and speakers, also resonate at certain frequencies, causing an audible hum. ANSI/IEEE standard 519 requires any device (whether intended as a communication device or not) that injects current on the grid to avoid doing so at certain frequencies and amplitudes to avoid causing flicker. Specifically, ANSI/IEEE standard 519 requires that no noise be added to the odd harmonics of the fundamental at or below the eleventh harmonic.

[0018] Despite the many engineering difficulties inherent in using the power grid as a communications medium, it has remained attractive to electrical utilities because the utility already owns the infrastructure, and it is present at all the points where the utility needs to collect data. Further, the regulatory and cost structure of publicly owned utilities (POUs) strongly favors them using owned assets (which can be profitably purchased and maintained via service rate increases) as opposed to paying operating expenses to a third-party communications provider such as a telephone or cable provider.

[0019] High-frequency transmissions (above 1MHz) are attractive because theoretically high data rates can be achieved. Such schemes, called BPL for Broadband over Power Lines, offer a potential theoretical bandwidth sufficient to deliver internet access to consumers via a gateway located in their electrical meter. In the early years of the twenty-first century, the Federal Communications Commission (FCC) in the U.S. actively promoted the concept of “Access BPL” as a means of delivering high-speed Internet access to rural American families. Long-range transmission of BPL signals, however, is impractical and expensive, because every transformer between the transmitter and the receiver must be fitted with a bypass or repeater mechanism, or the low-pass filtering characteristics of the transformer will block the signal. In the United States, where the number of consumers per service transformer tends to be very small – in rural areas often only one – the cost to implement BPL becomes prohibitive. Additionally, RF interference caused by BPL transmission has created opposition from aviation, commercial radio, amateur radio, and other sectors. The FCC has attempted to be supportive of BPL technology, but new compromise rules requiring BPL installations to

be capable of notching out (avoiding) frequencies where interference is reported have increased the complexity of managing a BPL service. Several attempts at deploying BPL consumer services have been abandoned.

[0020] Under the umbrella name Power Line Communication or PLC, some medium-frequency power line protocols have been used with success for Smart Grid applications, especially in Europe (and other locales with European-style grid architectures), where the number of consumers per service transformer is much larger than in the United States. The two most commonly used medium-frequency PLC initiatives are PRIME and G3, both promoted by commercial alliances based in Europe. PRIME uses orthogonal frequency-division multiplexing (OFDM) at the physical layer, with 512 differential phase-shift keyed channels (DPSK). PRIME achieves data rates as high as 128.6kbps, but is most reliable at 21.4kbps. Its frequency range is 42-89 kHz. G3 uses a similar physical-layer combination of OFDM and DPSK, offering 256 channels between 35 and 91k Hz with a data rate of 33.4kbps. Both G3 and PRIME are still sharply attenuated by transformers, though in most cases a receiver located on the medium-voltage side of a service transformer can successfully read meter transmissions from low-voltage sites served by that transformer, provided that the receiver is sited close enough to the service transformer. For these reasons, Smart Grid technologies based on these protocols are common in Europe and Asia. PLC protocols are also well-adapted to short-range power-line applications such as arbitrating the charging of electric vehicles.

[0021] At the other end of the spectrum are ultra-low frequency systems, chiefly used for control systems because they have little data-bearing capacity. Audio Frequency Ripple Control (AFRC) systems are used mostly in rural areas for load management: to turn off high-consumption devices such as electric heaters and air conditioners during peak load times, or to control use of other constrained resources, such as automated farm irrigation systems. An AFRC transmitter sits on the high-voltage side of a substation or transmission transformer and may service multiple substations. AFRC data rates vary from 2 to 10 bits per second, and the maximum message length is about 100 bits. After such a transmission, the transmitter requires a long idle period before it can transmit again, with a maximum duty cycle on the order of 10%. AFRC systems cause obvious flicker, but not at a dangerous frequency. Because they are typically used in areas of low

population density and transmissions are infrequent, the side effects are tolerated. AFRC is a broadcast technology operable from high voltage to low voltage and thus cannot be used for collecting meter data or other data about edge conditions, because that requires transmitting from low voltage to higher voltage.

[0022] Aclara®'s TWACS® technology operates by injecting pulses onto the power line when the fundamental power carrier crosses the zero point – twice per 50Hz or 60Hz cycle. This method operates either from substation to edge or edge to substation, and uses a polling protocol to avoid having one edge transmission interfere with another. It is slow because it is tied to the fundamental, and because of the polling architecture. It has been criticized by consumer groups for the amount of impulse and broadband noise it introduces onto the grid.

[0023] Landis+Gyr employs a low-cost, low-frequency edge transmitter originally developed by Hunt technologies, intended to operate in conjunction with AFRC to provide two-way communication over long distances on the grid. The data transmission method using this transmitter is cheap and reliable, but limited. It induces sympathetic current oscillations by connecting variable impedance to the power line. The data rate is low because the transmitter is dependent on a voltage relative to the power carrier, so that only a few pulses can be injected per 50Hz or 60Hz cycle. To achieve enough redundancy for detection at the receiver, the same signal must be repeated for several cycles, resulting in a data rate measurable in cycles per bit rather than bits per cycle. The method is also very noisy, in that each pulse resonates across a broad frequency band.

[0024] Despite their limitations, low-frequency systems such as those from Aclara and Landis+Gyr have achieved market penetration in rural areas where wireless systems are cost-prohibitive.

[0025] The problems with, and limitations of, the high, medium, and low-frequency PLC methods as discussed above have led in the 21st century to rapid development of custom built wireless networks for AMI data collection in the U.S. High-frequency on-grid methods have proven to be too expensive, not sufficiently reliable, and too fraught with error and uncertainty to be commercially viable. Low-frequency methods can be implemented with low-cost edge-to-substation transmitters, but these lack the data-bearing capacity required by modern AMI, and on-grid low-frequency substation-to-edge

transmitters like AFRC are large, expensive, and have undesirable side effects which limit their use in urban settings. One possible option would be to use high-frequency substation-to-edge transmitters in conjunction with low-frequency edge-to-substation transmitters. However, in the U.S. market forces have led to the rapid penetration of wireless AMI systems, especially in urban and suburban areas.

[0026] Cost constraints and availability of unregulated spectrum have dictated the use of mesh architectures at the edge of the AMI networks, with neighborhood concentrators that collect data from the meters and use traditional infrastructure (fiber or cellular) for backhaul to data centers. Mesh architecture means that although the RF transceivers used have individually high data rates, the edge networks are easily saturated. Most of the available data bearing capacity in these networks is used just for reporting meter interval data, with limited capacity reserved for firmware updates and control packets for applications such as demand management.

[0027] There are two major factors that limit the utility of the existing AMI infrastructures. The first is, of course, the capacity limitations of the mesh. The second, and more significant, is the fact that the AMI network is not congruent with the electrical grid. It is capable of providing little information about the operational state of the grid. This is unnecessary for meter reading, but more sophisticated Smart Grid applications for energy conservation, asset protection, load balancing, fault isolation, and recovery management require accurate information about the schematic relationship of grid assets, load and conditions on the several segments of the grid, and the current state of bi-modal and multi-modal assets. This information, together with the geospatial locations of the same assets, is called the Grid Map.

[0028] Utilities typically maintain two maps or models of the Grid. A Physical Network Model (PNM) aggregates the geospatial location of the assets on the grid. PNMs, thanks to modern GPS technology, are reasonably accurate with respect to point assets such as substations, capacitor banks, transformers, and even individual meters. Inaccuracies stem from failure to update the maps when repairs or changes are made. For example, a service transformer may move from one side of a street to the other as a result of street widening. Longitudinal assets, especially buried cables, are less well represented in the PNM. The PNM can contain as-designed data, but since in many places the cable was

laid before global positioning technology had matured, the designs were based on ground-level survey, and the original maps may or may not have been updated to reflect changes. Subsequent surface changes complicate the problem of verifying the geographic path taken by medium-voltage distribution lines.

[0029] The second model is the Logical Network Model, or LNM. LNMs describe how grid components are connected, without reference to their geospatial location. The LNM changes frequently. During the course of repairs, the way transformers attach to taps and laterals, and meters attach to transformers, may be altered. Such changes affect both the LNM and the PNM. In many utilities, such changes are recorded manually by field agents. The manual reports may or may not be updated in the LNM and PNM, and when updates are made the time lag between maintenance occurring and its being recorded is variable. Additionally, many grid components, especially regulators, switches and reclosers, change state automatically. Unless these components are instrumented with communications back to a data center rather than simply being subject to local control systems, such dynamic changes are not reflected in the LNM. They do, however, affect the power path, the load and environmental stress on other components of the distribution grid, and the service level to consumers.

[0030] Examples of significant but not reliably known aspects of the (actual) Grid Map are the feeder and phase by which each meter is currently powered, the relative load on each phase of each feeder, especially on subordinate branches (laterals) of the grid, the actual voltage supplied to each meter, the power factor along the edges of the grid, whether all the power drawn at a transformer is metered, and the state of switch sets, especially after a weather event that has caused outages. If this information were reliably known, utilities could conserve energy, much of the savings from which would pass on to consumers, save maintenance costs, prolong the life of equipment in the field, improve the efficiency and life of utility and consumer equipment, avoid outages, and reduce recovery times after unavoidable outages.

[0031] The problem of automated, dynamic grid mapping is not solved by wireless Smart Meters. Smart meters can measure and record current, voltage and power factor (or phase angle) at the meter, but because they have limitations on how much data they can store and how much data capacity is available for transmission, utilities may choose not

to program the meters to report these data. The other data elements described cannot be detected by most modern AMI systems. U.S Patent 7,948,255 to Coolidge, et al. discloses instruments for phase detection. However, the instruments in Coolidge are intended to be used by field engineers rather than incorporated into the Smart Grid.

[0032] The consensus among utilities is that the volatility of the LNM is such that using field engineers to measure and monitor changing attributes of the grid map is not a cost effective or workable solution. For example, conservation voltage regulation efforts were undertaken in the 1990s based on static measurements, and subsequently abandoned because the measurements became outdated too quickly. Today, utilities habitually oversupply consumers, delivering an average effective voltage of 122vAC to a 15 or 20 amp-rated circuit in a residence to ensure that fluctuations in load, power losses in the home wiring, etc. do not result in some consumers' service falling below 110vAC effective at individual outlets inside the building, which is generally the optimum for home appliances and the like. The goal of a well-instrumented fine-grained conservation voltage regulation system might be to reduce the typical effective voltage at a single-phase meter to 114vAC as measured from one leg of the typical 240vAC service to neutral. 114vAC effective at the meter is as low as it is reasonable to go without risking under-powering some outlets in the building, (i.e. not less than 110vAC at any outlet) due to additional losses which are typical inside the home or office.

[0033] Since electrical devices consume more energy when powered at the high end of their rated range, this practice of over-delivering impacts consumers' electric bills directly, as well as forcing generation-poor utilities to buy power, increasing their costs. Ultimately, the practice results in more fossil fuel being consumed than necessary.

[0034] Cost constraints also dictate that placing SCADA instrumentation at every medium-voltage field asset is impractical. The "touch points" on the distribution grid are, for better or worse, largely the meters at the edge and the instrumentation in the substations. This dictates that techniques for power line communication be revisited, because signals traveling on the power line can be used both to infer and to report grid mapping information not detectable by means of wireless AMI. The ubiquitous presence of wireless AMI for reporting meter data can be considered as a benefit in the search for effective grid-mapping technology, in that it frees the limited data-bearing capacity of

low-frequency on-grid transmission methods to support grid mapping systems instead. It is, however, needful to identify a transmission method that is low cost at the edge, coexists with the AMR or AMI, and does not trigger any of the above-noted pitfalls of on-grid transmission: requirements for intermediate devices such as repeaters between the edge and the substation; unacceptable flicker; RF interference; impulse noise; etc. Finally, the transmission must require very little power, because the energy expended driving the transmitters reduces the energy conservation benefits obtained.

[0035] As discussed above, some existing PLC methods have adapted radio modulation techniques and channel access methods to the medium of the electrical distribution grid. For example, PRIME uses FDMA with DPSK.

[0036] In addition, Code Division Multiple Access (CDMA) is a channel access method most famously used in cellular telephony standards cdmaOne, WCDMA, and CDMA2000. CDMA spreads its signal across a range or *band* of frequencies, as do other similar technologies; hence the term *broadband*. Multiple access means that more than one transmitter can use the same channel (here, a power line) without the signal from one transmitter destructively interfering with the signal of another transmitter. In CDMA, each transmitter which uses the same band is assigned a distinct reference code or *chip*. The transmitted signal equals the exclusive OR (XOR) of the chip with the data signal. If the chips (treated as binary vectors) are mathematically orthogonal, then the receiver can separate out the several data signals from the additive received waveform. A requirement of standard CDMA as used in a wireless application is that there is a dynamic feedback loop from receiver to transmitter to ensure that the power of the several signals received from the different transmitters is the same or nearly the same at the receiver. The feedback loop permits the transmitters to rapidly and dynamically adjust their transmission power to maintain the balance.

[0037] Frequency Division Multiple Access (FDMA) means that multiple channels in a medium are created by having different transmitters use different frequencies (or different frequency bands). A signal injected on the power line creates harmonic signals of different amplitudes. If the frequency-division bands are incorrectly chosen, then the harmonics from different bands can coincide and create false signals that interfere destructively with the intended signals. The obvious means of eliminating this effect is to

place the channels far apart on the frequency spectrum. This, however, reduces the overall data bearing capability of the medium by “wasting” spectrum.

[0038] A third channel access method is Time Division Multiple Access, or TDMA. In TDMA, the channel is divided cyclically over time, with each transmitter sharing the channel assigned a specific time slot in the cycle where that transmitter uniquely has permission to transmit. TDMA requires that all the transmitters have system clocks which are synchronized with one another within a close enough tolerance that one channel accessor does not overlap its transmission with that of another channel accessor.

SUMMARY OF THE INVENTION

[0039] The present invention includes a system comprising at least one intelligent edge transmitter called a Remote Hub Edge Transmitter, each small enough to reside inside a Smart Meter. A Smart Meter that contains a Remote Hub Edge Transmitter is called a Remote Hub GLA Smart Meter, or simply a Remote Hub. The Remote Hub transmits by injecting a modulated current into a power main that supplies an electric meter. The system also includes at least one receiver sited at at least one electrical distribution substation operable to receive transmissions from the intelligent transmitters. No additions or alterations to the distribution grid between the Smart Meters and the substation are required to allow the receiver to reliably detect and decode transmissions from the edge transmitters. The system further comprises one Computing Platform for each substation that contains at least one receiver, the Computing Platform having access to a conventional high speed network such as the Internet for transmitting data acquired from the at least one receivers to a data center at which the received data is used by a Concentrating Computer System, or Concentrator, to update other utility systems such as, but not limited to, the LNM and PNM. In some microgrid deployments the Computing Platform and the Concentrator may be the same server, with the data center sited inside the service area of the microgrid. The system may additionally comprise Smart Meters or other devices, such as field-deployed switches and voltage regulators, which are not Remote Hubs, augmented by intelligent platforms operable to employ short-range PLC transmissions using a well-known protocol such as G3 or PRIME to communicate with at least one Remote Hub. Such augmented devices which are not Remote Hubs are designated as Subordinate Remotes, and any augmented device, without regard for

whether it is a Subordinate Remote or a Remote Hub, may be referred to generally as a Remote. Each Remote Hub manages only Remotes powered by the same service transformer as the Remote Hub. A short range on-grid network consisting of a collection of Remotes comprising at least one Remote Hub and zero or more Subordinate Remotes is called a Transformer Area Network or TAN. A Remote is any device that can communicate in a TAN. Subordinate Remotes and Remote Hubs are classes of Remotes. All Remotes have TAN transceivers and all Remote Hubs have Edge Transmitters.

[0040] Subordinate Remotes, Remote Hubs, the Substation Receiver, and the associated Computing Platform and Concentrator all contain stored programs on non-volatile computer-readable memory on which are stored instructions for operating a Grid Location Aware™ (GLA) network. The Subordinate Remotes, Remote Hubs, the Substation Receiver, and the associated Computing Platform and Concentrator also contain processing units (CPUs) that execute the stored instructions allowing each node in the network to implement methods for organizing the on-grid network and transmitting and receiving messages on the network in order to permit other methods embodied as stored programs and executing on the at least one Substation Receiver, Computing Platform and Concentrator to detect and infer schematic grid location attributes of the network and publish the detected and inferred attributes to other application systems including geospatial information systems maintaining the logical and physical network model.

[0041] One method implemented by the Remote Hubs and the Substation Receiver provides for channelizing and modulating current signals transmitted from the at least one Remote Hub in the service area of an electrical distribution substation such that the signals are received at the Substation Receiver and the Substation Receiver is able to infer the electrical phase of the specific feeder upon which the signal was transmitted. The signal is transmitted on a broad band of the frequency spectrum called a channel, but the frequency bands of channels are selected so that the frequency is lower than the low-pass threshold of the service transformer that powers the Edge Transmitter. Several modulation techniques have been used in this context, including frequency spread modulation, Binary Phase-Shift Keying (BPSK), and Quadrature Phase-Shift Keying (QPSK). Higher-order modes of phase-shift keying (mPSK) may be used. However,

BPSK and QPSK may be preferred embodiments along with frequency spreading, because higher-order PSKs require more power at the transmitter in order to achieve the same signal strength at the receiver. According to some embodiments of the methods, an Edge Transmitter is capable of encoding at least 80 bits per second of post-FEC (forward error correction) data in a burst transmission at low but adequate current so that the signal is not so significantly attenuated by intermediate transformers, capacitors, long lines, underground wiring, and the like to prevent reception by the Substation Receiver. In other embodiments, an Edge Transmitter may be capable of encoding at lower bit rates. Encoding at lower bit rates improves reliability, but limits the amount of data transmitted. In order to obtain the same post-FEC message success rate while transmitting at at least 80 bps, different modulation types may require different Forward Error Correction rates. The method requires little power to inject the signal, so that the signals as modulated do not radiate energy in the RF spectrum or cause flicker or hum on devices in proximity to the transmissions or exhibit any of the other undesirable characteristics of prior art methods of on-grid messaging. The method works on all the grid topologies described herein above, and can support a sufficient number of Remote Hubs per substation transformer that even the largest substations can be fully covered by the resulting Grid Location Aware™ network.

[0042] The Substation Receiver may also implement a variety of methods of sampling the ambient waveforms at a multiplicity of frequency bands on the power lines, filtering out the high-energy harmonics of the fundamental power wave, detecting the signal on one or more of a plurality of power lines (comprising each of the three phase lines of each feeder emanating from each bus of a given substation transformer), inferring the phase and feeder combination on which the signal was transmitted based upon a comparative analysis of each of the power lines, ranking them based on the signal quality, error performance, and/or amplitude versus frequency at a multiplicity of points throughout the spectrum of interest. When the Substation Receiver has completely processed a transmission, it packages the decoded transmission together with any additional information about the message inferred by the receiver logic, such as the phase and feeder on which the message was transmitted, the channel on which the message was transmitted, and an indication of the parameters of the modulation method used for that

transmission. The Substation Receiver forwards the entire message package to the substation Computing Platform using a normal TCP/IP-based protocol such as HTTP.

[0043] Another aspect of the invention is a method of identifying which frequency bands are the best data carriers at each substation transformer, defining at least one data-bearing channel on the candidate frequency bands, optionally defining a series of time slots on each channel in which edge devices may transmit, selecting a modulation technique, and, if frequency spreading is the chosen modulation technique, defining a set of at least one orthogonal codes or “chips” per channel to be used for modulating transmissions. The combined channelization model is then employed by the method to provision the collection of GLA Smart Meters, including both Remote Hubs and Subordinate Remotes, supplied with power by a substation to assign to each of the at least one Remote Hubs a policy describing on which frequency-based data-bearing channel(s) the Remote Hub may transmit, and under what circumstances the Remote Hub must transmit. The policy describes multiple aspects of the channels, including modulation method, frequency bands, chip selection algorithm if chips are used, and message preamble pattern. Frequency-based channels and chips must be assigned in such a way that transmissions are not destructive when segments of the grid are, for example, switched from one substation transformer to another. The provisioning scheme anticipates and minimizes the problem of crosstalk, and provides means for logic on the Substation Receivers, the substation Computing Platforms and the Concentrator to hierarchically process the messages received from each Remote Hub and use them to infer the state of non-edge features of the grid, such as switches, reclosers, and breaks in the power lines. Other properties of the transmission are determined dynamically by firmware and instrumentation on the Remote Hub. For example, the power used when transmitting may be related to the impedance of the line as measured immediately prior to transmitting.

[0044] In some embodiments of the invention, a number of techniques may be employed for managing channel quality, depending on the availability of Substation-to-Edge broadcast capability from adjacent networks, such as an AMI, AMR, and/or radio broadcast transmitter. Software on the Substation Receivers and Computing Platform may monitor aspects of the channel quality and take measures to ensure that messages

from the Remote Hubs experience an acceptably high success rate. According to one aspect of the invention, an acceptably high success rate may be ensured by rotating the responsibilities of the several channels, except that at least one non-structured channel is not rotated but remains dedicated to provisioning and alerting. For example, if two data bearing channels have been identified, and one data bearing channel demonstrates a higher success rate than the other, then the network may be provisioned to have Remote Hubs alternate between transmitting on the better channel and the other channel. This reduces the overall probability of a given Remote Hub experiencing an unacceptably high message failure rate.

[0045] Other options for channel management may be to alter the definition of a channel so that the channel has a wider frequency spread, and/or uses more FEC bits per burst. Still another option is to move a channel to a different place in the spectrum, either permanently, or at different times of the day based on an observed cycle in impedance, impulse noise, or some other characteristic of the channel relevant to message success rate. None of these mechanisms require a fast feedback loop between the Edge Transmitters and the substations, as is the rule with some modulation techniques such as CDMA. Rather, the apparatus at the substation conducts a time-duration analysis of the behavior of the network, and then broadcasts new provisioning policy based on the analysis. Many characteristics of the network may be taken into account when making a policy change, such as observed patterns of crosstalk, variations in impedance, harmonic mixing, and the like. A policy change may impact multiple substations, which may be interconnected by switching systems or other forms of redundancy.

[0046] Yet another aspect of the invention is a method employed by the stored programs at each of the at least one Remote Hubs to integrate the Edge-to-Substation GLA network with adjacent networks, such as the AMI (regardless of the AMI architecture) and the higher-frequency PLC based Transformer Area Network, as well as with the native intelligence of the Smart Meter platforms themselves. In this method, the Remote Hub, whose high-frequency PLC protocol stack (e.g. PRIME) enables it to act as the master node in the TAN, carries out the TAN-management activities. TAN management activities include, but are not limited to, polling the PLC protocol stack to detect newly discovered Subordinate Remotes. The Remote Hub also polls the local native Smart

Meter intelligence to obtain local data such as current, voltage, and phase angle, and polls the reachable population of Subordinate Remotes to obtain similar data from the native Smart Meter intelligence at the Subordinate Remotes. The Remote Hub then derives the results from the collected data points. The Remote Hub stores, compresses, and/or processes the gathered data according to the policies and application algorithms on the Remote Hub until the operable policy dictates that the gathered data and/or derived results of the gathered data may be transmitted to the Substation Receiver by the Edge Transmitter module of the Remote Hub. By way of example, the Remote Hub might sample RMS voltages from each Remote including itself for a period of one hour, storing and/or summing the samples, then compose a message containing the highest and lowest voltage measured, the identity of the Remotes where these outlying values were measured, and the average voltage over all the Remotes for the interval. Such “derived results” are useful in providing a characterization of conditions in the TAN without requiring that a large amount of data be transmitted to the substation. These policy-based “derived results” can be thought as “apps” for Remotes and can be altered or changed from time to time.

[0047] The Remote Hub is further responsible for using its provisioned policy and discovered TAN configuration data to determine when it is appropriate to format, encode, and transmit an alert message on an edge-to-substation channel. Such messages may include pairing messages that report the discovery of a new Subordinate Remote; pairing alerts that report the loss of communication with a known Subordinate Remote; other alerts that report changes in the TAN or at meters (such as power surges, sags, and spikes); and scheduled data reports that transmit the data collected and derived from the native Smart Meter intelligence in the TAN to the apparatus at the substation. In some embodiments of the invention, channels are not time-slotted, and Remote Hubs may transmit only exception reports or computed data reports on a randomized posting schedule in which an adequate number of transmissions are performed to provide an acceptable probability of achieving at least one successful transmission at the desired rate. All conditions that ultimately trigger an Edge Transmitter message, from a detected failure in a Remote to achieving the conditions required for computing and reporting a derived result, can be described as events.

[0048] If slotted channels and/or time-scheduled transmission policies are used, the Remote Hub may require a method of synchronizing its system clock to a known tolerance with other Remote Hubs in the same service area. Each Remote Hub may poll the local meter or AMI network to obtain the AMI network time that the Remote Hub uses to determine when scheduled transmissions must occur, and to obtain data blocks received via the AMI that are intended for the Grid Location Aware™ intelligence on the Remote Hub or on the Subordinate Hubs. Such data blocks may include firmware updates and changes in network policy or provisioning that will affect the subsequent behavior of the Subordinate Remote. The Remote Hub distributes firmware updates and policy changes to the Subordinate Remotes as necessary via the local PLC channel of the TAN. Additionally, Remote Hubs may synchronize based on a wireless broadcast signal. If no synchronization method is available, channel access may not be based on time slots at all. This reduces the data-bearing capability of the network but does not impact the ability of the system to provide grid-location data. In some embodiments, Remote Hubs and/or Subordinate Remotes may contain a Global Positioning System (GPS) receiver. The GPS signal may be used for synchronizing the Remote Hubs in addition to providing means to associate the logical network model with the physical network model.

[0049] In still another aspect of the invention, the Computing Platforms and the Concentrator maintain two master data tables that can be initially extracted from the utility's PNM and/or LNM, or which can be entirely accumulated from reports from the Remote Hubs. These data tables are the Inventory, which is a table of all the Remote Hubs and Subordinate Remotes that have been detected, and the Grid Map, which is a schematic representation of the grid's topology and state, similar to an LNM. The Grid Map and Inventory at substation Computing Platforms may be partial, representing only the portion of the grid accessible to the substation at least at certain times. The Grid Map and Inventory at the Concentrator generally represent the entire utility service area, although gaps in the Grid Map may exist if instrumentation of the service area with Remote Hubs and Subordinate Remotes is incomplete. When the Computing Platform at a substation receives any message from a Subordinate Remote, it compares the data in the message and the message enhancements inferred by the Substation Receiver with the data in the Inventory and Grid Map. The logic and policy on the Computing Platform are

used to determine if the local copy of the Grid Map and Inventory need to be updated, and whether the update must be sent on to the Concentrator to update the master Grid Map and Inventory. If the policy in effect at the Computing Platform so dictates, the data collected from the edge is also forwarded to the Concentrator. The Concentrator in turn carries out policies dictating which events and scheduled reports must be published out to other data center applications. Additionally, a transactional history of each instance of detecting a transmission on an Edge-to-Substation channel may be maintained as a data archive, including not only the message of the transmission, but other observed and measured properties of the transmission that may be useful in analyzing and optimizing the performance and utility of the on-grid data collection network. The data archive of the transaction history may be maintained in the Computing Platform, or it may be uploaded to a Concentrator. However, the primary use is local to each Computing Platform, where analysis of local grid characteristics over time may be used to provide Remote Hubs with data and policies describing how to maximize the probability that their transmissions are received correctly.

BRIEF DESCRIPTION OF THE DRAWINGS

[0050] The accompanying drawings, which are incorporated in and form a part of the specification, illustrate the embodiments of the present invention and, together with the description, serve to explain the principles of the invention.

[0051] Figure 1 is a simplified illustration of the power path from a generation point to a distribution substation to a consumer, showing the high voltage, medium voltage, and low voltage regions of the distribution grid and depicting some of the major features of an electrical distribution grid.

[0052] Figure 2a is a simplified fragment of a radial-architecture distribution grid showing the lack of cycles in the grid topology.

[0053] Figure 2b is a simplified fragment of a looped-architecture distribution grid depicting two substations each able to deliver power to the service transformer delivering low-voltage power to the group of residences shown. The substation at left is currently powering the residential group.

[0054] Figure 2c is a simplified fragment of a networked architecture distribution grid. The four feeders shown could originate at a single substation (typical) or at multiple

substations. The rectangular grid connects service transformers peer-to-peer on the low-voltage side so that all feeders deliver power to the loads below the substations concurrently.

[0055] Figure 2d is an exemplary simplified fragment of a campus network. A three-phase transformer powers a 480 Volt Bus from which depend a number of three-phase laterals which run through the campus powering individual electrical outlets. Adding low-voltage generation points to the bus and providing means to isolate the bus from the distribution line converts the campus network into a self-sufficient microgrid.

[0056] Figure 3 is a high-level software deployment model of a Grid Location Aware™ network including back office features, substation apparatus, and transformer area networks (one expanded) including intelligence at remote hubs and subordinate remotes.

[0057] Figure 4 is a simplified block diagram of the substation apparatus in a Grid Location Aware™ network, illustrating how the Grid Location Aware™ network apparatus couples to the existing SCADA lines in the substation and how the data from the Substation Receiver is backhauled to a data center.

[0058] Figure 5 is a block diagram of the multi-threaded software architecture in the Substation Receiver showing how Edge-to-Substation signals are acquired, channelized, detected, demodulated, decoded, and stored for processing and backhaul.

[0059] Figure 6 is an elevation of a Remote Hub GLA Smart Meter.

[0060] Figure 7a is a top view of an accurate 3-dimensional model of one embodiment of the Edge Transmitter module of the Remote Hub GLA Smart Meter for a Form 2S residential meter.

[0061] Figure 7b is a bottom view of the same model of the Edge Transmitter module.

[0062] Figure 7c is a schematic block diagram of the electronic components of the Edge Transmitter module of the Remote Hub GLA.

[0063] Figure 7d is a detail of the isolation circuit of the Remote Hub.

[0064] Figure 8 is an elevation of a Subordinate Remote GLA Smart Meter.

[0065] Figure 9 is a schematic block diagram of the electronic components of the PLC communication module of a Subordinate Remote.

[0066] Figure 10a is a graph (not to scale) of a snapshot of the AC waveforms on a distribution grid at a meter connection point. The power fundamental and its odd

harmonics are highlighted, and three CDMA-like broadband frequency-divided Edge-to-Substation channels are shown, one extended over time to illustrate time divisions.

[0067] Figure 10b provides a time-domain view of three Edge-to-Substation channels showing two scheduled and one unscheduled channel.

[0068] Figure 10c illustrates a typical structure of a single Edge-to-Substation message burst in one embodiment of the invention.

[0069] Figure 10d illustrates an alternative structure of an Edge-to-Substation message burst in another embodiment of the invention.

[0070] Figure 10e illustrates how using multiple chips on the same frequency band may prevent collisions.

[0071] Figure 11 illustrates communication paths local to the Transformer Area Network, both within a GLA Smart Meter and between the Remote Hub and a Subordinate Hub. Note that in Figure 11 the elevations of the Remotes depict embodiments where a secondary communication module is not employed.

[0072] Figure 12a illustrates a three-phase service transformer with three TANs.

[0073] Figure 12b illustrates a three-phase service transformer with one TAN and Proxy Remotes.

[0074] Figure 13 illustrates how the usage of frequency bands available for Edge-to-Substation transmissions may be manipulated in order to optimize the utilization of the usable frequency for best network performance.

DETAILED DESCRIPTION OF THE INVENTION

[0075] The present invention comprises a system and methods for constructing and operating an on-grid data collection network in such a way as to integrate the network with other adjacent networks and devices present at the edge, substations, and features of an electrical distribution network, wherein the other networks and devices may include Smart Meters and the AMI and a conventional network such as the Internet. The system and methods further integrate the data collected by the on-grid data collection network at a data center and may publish the collected data to other applications. The system and methods may also employ the capabilities of the integrated networks to infer otherwise unknown static and transient attributes of the electrical distribution grid and report them

via the integrated networks for the purpose of improving the physical and logical network models of the Smart Grid. This leads to the ability of the models to support Smart Grid applications such as conservation voltage reduction, volt/Var optimization, load balancing, fault isolation, and recovery management.

[0076] Figure 3 is a logical block diagram of the intelligent platforms of one embodiment of the invention showing where the different intelligence resides with respect to a converged network comprising a conventional wide area network such as the Internet, an Advanced Metering Infrastructure, the medium voltage electrical distribution grid, and the at least one Transformer Area Network connected at the edge of the medium voltage electrical distribution grid via at least one service transformer. Figure 3 shows that the system may be divided into three regions or *tiers*. The edge tier **301** comprises at least one Transformer Area Network (TAN) **302**. Each TAN comprises a Service Transformer **303**, at least one Remote Hub **304**, and zero or more Subordinate Hubs **305**. When Subordinate Hubs **305** are present, the Remote Hub communicates with the Subordinate Hubs via a standard short range, PLC protocol, such as PRIME. The Remote Hubs **305** may be operable to transmit current-modulated signals. To avoid having multiple Remote Hubs unintentionally installed on a TAN, the installation procedure may provide a mechanism to allow a newly installed Remote of either type to detect whether Remote Hub is already present on the local TAN. The invention does not require the installation of communication devices or other modifications between the edge tier and the substation tier. However, if it is desirable to collect data from a feature of the medium-voltage grid such as a capacitor bank, a variation of the Remote Hub device may be installed there. Such a Remote Hub is technically still at the Edge Tier, since it is powered by a low-voltage outlet located at the feature, and not directly from the medium-voltage line(s) upon which the grid feature is installed.

[0077] Still referring to Figure 3, the substation tier **306** comprises at least one Substation Receiver **307** operable to receive transmissions from the Remote Hubs in the edge tier without recourse to any signal amplifiers, bypass mechanisms, or bridges installed on the medium voltage infrastructure of the electrical distribution grid. The Substation Receiver(s) connect via a local area network to a Computing Platform **308** containing non-volatile computer-readable memory and a CPU for storing and executing the

software **309** which maintains the Inventory and Grid Map databases and carries out the tasks of provisioning and managing the converged data network. Additionally the Computing Platform stores and executes software **310** which processes the Inventory and Grid Map data in combination with messages received from the Substation Receiver **307** to infer information about the state of the Grid over and above what the Substation Receiver alone can detect based only on incoming transmissions. In some embodiments of the invention, the computer-based components of the Substation Receiver and the Computing Platform components are hosted on the same server. In such embodiments, the communications protocol (such as HTTP) used to transfer data between the Substation Receiver and the Computing Platform software components need not change, even though there is no physical local area network required. Computing Platform **308** connects to a conventional wide area network **311**, such as the Internet, for the purpose of communicating with a Concentrator **312** in the data center tier **313**. In some embodiments of the invention, and regardless of whether the Computing Platform and Substation Receiver are the same server or separate servers, the servers may be configured in a redundant cluster to ensure continuous operation of the system.

[0078] Referring again to Figure 3, the Concentrator **312** hosts software with an analogous architecture to the software in the substation(s), comprising a network and data management component **314** providing software services to one or more applications **315** for Grid Location Awareness (GLA). The applications use conventional network-based messaging protocols such, as but not limited to, JMS, SOAP, and REST to publish information to subscriber applications such as a Geospatial Information System (GIS) **316**. The data and network management component **314** may integrate with AMI head-end **317** for the purpose of causing the AMI network to broadcast data blocks to the Remote Hubs in the edge tier **301**. Data and network management component **314** may integrate with AMI head end **317** using a standard protocol and/or a proprietary interface defined by the AMI vendor.

[0079] Other embodiments of the invention may include the convergence of alternative ancillary network components. For example, Substation-to-Edge broadcast capability and/or time synchronization from the substations to the Remote Hubs may be provided by medium-voltage PLC transmitters attached to the feeders at the substation rather than

using an AMI for this purpose. Likewise, a separate radio transmitter broadcasting messages originating at the substation may be employed. The radio transmitter does not need to be physically located at the substation as long as there is a low-latency network connection from the Computing Platform at the substation to the transmitter. The same radio transmitter may serve as the Substation-to-Edge channel for a multiplicity of substations. When the Substation-to-Edge channel is not an AMI, synchronization of the Remote Hub clocks may be provided as described in US Patent Application No. 13/566,481, titled System and Methods for Synchronizing Edge Devices on Channels without Carrier Sense, which is incorporated herein by reference. In embodiments of the invention where channels are not slotted, clock synchronization is unnecessary.

[0080] Figure 4 details how in one embodiment the Substation Receiver **401**, here shown co-hosted on a single server with the other software components of the Computing Platform, monitors the feeder lines **402** on the low-voltage side of the substation transformer **403** by attaching secondary current sensors **404** to the SCADA loops **405** already in place. The secondary current sensors provide inputs to the Substation Receiver. This coupling method allows a Substation Receiver to be installed on a substation transformer without disrupting the operation of the substation. Other coupling methods such as hot-stick clamp-on current transformers are well known in the art, and may be equivalently employed in lieu of the secondary coupling to SCADA loops method described herein. Some substations may lack SCADA loops, or they may be inaccessible due to physical placement or due to utility regulations.

[0081] Figure 5 details the software architecture and method used by one embodiment of Substation Receiver logic **501** to capture, detect, differentiate, and decode the multiplicity of signals being received from the Remote Hubs at the edge of the portion of the distribution grid supplied by the substation transformer associated with this Substation Receiver. Inputs from the GLA CT lines **504** arrive at data acquisition module (DAQ) **502** in real-time as raw digitized signals where they are then buffered and recorded on ramdisk **503**. Operating in parallel with the DAQ process, Channelizer **505** reads the raw signals and reorganizes them by frequency band of interest into channelized signals stored on ramdisk **506**. Operating in parallel with the Channelizer, Preamble Detector **507** samples the channelized signals received on every feeder-phase attempting to

recognize the one or more preamble patterns which precede every transmission. In one embodiment, the Preamble Detector looks for all legitimate preambles, thereby allowing for the receipt of transmissions that are outside their allocated time slots or which were transmitted on a non-slotted channel. The Preamble Detector may also use its knowledge of the channel time slots in order to minimize the amount of channelized recordings it must sample. In one embodiment, when the Preamble Detector finds a preamble, it determines and marks the point or points in the data stream at which the Demodulator **509** should start processing. Multiple copies of the same message may be detected due to crosstalk. All copies are retained for demodulation. Preamble Detector **507** provides the marker information and Channelized signal data to the Demodulator through ramdisk **508**. Operating in parallel with the Preamble Detector, Demodulator **509** reads all copies of the messages from ramdisk **506**, uses the frequency bands and possibly time slots in which the messages were found and the known policies of the Remote Hubs to determine how to decode the message. Policy elements may include the modulation technique in use on the channel, the time based access policy in use, and, if frequency-spread modulation is used, which chips could have been used in modulation. Sometimes there may be more than one possible choice of chip. If the demodulator attempts to apply the wrong chip, this will be indicated by a demodulation error and/or FEC failure. This parallelized embodiment of the receiver logic permits multiple processor cores to operate on each message stream, with the modules early in the data flow operating on later transmissions, while modules late in the data flow operate on earlier transmissions.

[0082] Still referring to Figure 5, Data Manager **514** may be responsible for synchronizing the several processes' access to the data stored on the ramdisks **503**, **506**, **508**, and **510**, so that each process at a given time is operating on completed data that was output by its predecessor process rather than attempting to access stale or volatile buffers. The Data Manager also can copy data from the ramdisks to a large archive disk **513** for later retrieval, study and post-processing. The archive disk **513** may store archived data about the history of at least one Edge-to-Substation channel, including each transmission as detected on however many channels upon which the transmission could be detected, together with potentially all of the measurements taken by the receiver during the detection and demodulation process, described herein below. Some embodiments of the

invention may archive all of the measurements and all copies of the message, while other embodiments may elect to archive only those measurements that support the types of post-processing that will be used at a particular installation. For example, some metrics are useful only where a Substation-to-Edge channel is present. Other metrics are not needed if only one Edge-to-Substation channel is employed. By the time the messages have been demodulated and written to ramdisk **510**, they have been enhanced with enough information to identify the transmitter and infer the feeder and phase on which the message was transmitted, as described in more detail below. The message bundle for each feeder-phase on which the message was received may include, but is not limited to, the signal amplitude at a multiplicity of characteristic frequencies within the full frequency band available for edge-to-substation transmissions, the signal quality as determined by Demodulator figure of merit measurements such as E_b/N_0 (the ratio of energy per bit to noise power spectral density) and message detection rate (if the preamble pattern is recognized, whether the associated message can be decoded), the time the signal was received, and the chip, if any, with which the message was encoded. Other examples of metrics that may be archived are bit error rate by channel, Gaussian noise levels and impulse noise events, symbol constellation separation, inter-symbol interference, receive time drift, and drift in modulation frequencies. These measurements both support the Grid Location inference and, when analyzed over time, contribute to the optimization of the network. For example, comparing the signal amplitude at multiple frequencies from one instance of a detected message against another instance detected at the same time helps to determine which instance is due to crosstalk and which instance is the original transmission. Finding trends in the amplitude over time, on the other hand, may point up the need to send feedback to a Remote Hub, as is described herein below. The resulting message bundles are passed to the Network Management and Grid Location Awareness software via interface **511**. These software components, not shown in Figure 5, but shown as **309** and **310** respectively in Figure 3, use the provisioning policies of the edge transmitters together with the message properties and the message content to determine which Remote Hub edge transmitter sent the message, and compare the signal characteristics of the copies of the message received on each different feeder-phase input where it was detectable in order to establish on which feeder and phase the

message was actually transmitted. This conclusion is compared with the information in the grid map to determine if a change in grid topology or state has occurred. This allows the Grid Location Awareness (GLA) algorithm to infer not only the phase of a meter where the phase was previously unknown, but also changes in switch states in loop or networked configurations and schematic alterations in any type of grid resulting from repairs and maintenance. How these measurements vary over time for each specific transmission source (Remote Hub) and for the several channels can be analyzed to determine what variations in transmission configuration and policy might improve the message success rate for individual Remote Hubs and entire channels. Consequently, the Grid Location Awareness software may archive the measurements and the grid location determinations made as a transaction history. These data may be retrieved and analyzed at a later time for the purpose of network performance optimization. Some measurements, including systematic drift in modulation frequencies, may become indications that a Remote Hub is malfunctioning and needs to be replaced. Other trends may lead to the substitution of one frequency-based channel for another, or to altering the times of day when certain Remote Hubs are directed by policy to transmit. When post-processing the archived data leads to recommendations of this kind, the recommendations may be forwarded to the Concentrator to be passed on to other applications (such as Service and Work Order management) or transmitted to Remote Hubs via a Substation-to-Edge channel. Additionally, relative signal levels measured at the Substation Receiver can be used to calculate equalization parameters for a Remote Hub. When communicated to a Remote Hub this allows the Remote Hub to pre-equalize (i.e. empathize certain frequencies) the transmitted signal. This can improve the message success rate.

[0083] Referring once again to Figure 3, the software components **309** and **310** on the Computing Platform **308** decode the demodulated, error-corrected message received from the Substation Receiver at the semantic level. The semantic decoding may include decryption and a CRC check on the decrypted message. This helps preclude the introduction of false data, for example due to tampering with the firmware on the Remote Hubs, or installation of a meter from a different service area on the network for malicious purposes. Once the message has passed this level of decoding, the data payload in the message may yield additional grid awareness information. A pairing message indicates

that a new meter has been installed successfully, or that a known meter is now connected to a different service transformer or a different phase of a multi-phase transformer. Scheduled data messages may provide information about voltage levels, demand, and power factor at the edge, as well as any other data or results computed from data available from the instruments at the Remotes, limited only by the data-bearing capacity of the channels. Even the failure of a scheduled message to arrive is informative, indicating that there may be an outage. Any copy of the message may be used to extract data, not only the copy from the feeder and phase on which the message was determined to have been transmitted. Sometimes the main copy may contain bit errors while crossover copies do not. These determinations may also be relevant metrics for performance optimization, and may be included in the archive of transmission-related metrics described herein above. The data archive may be indexed so as to facilitate aggregating historical results by channel, by individual transmission point, by time of day, by configuration parameters and policies in effect at the time of transmission, and other factors that may be relevant to performance outcomes such as message success rate and the confidence level of the inferred grid location.

[0084] When all the information has been extracted from a message bundle at the substation, the software components on the Computing Platform apply a policy to determine what data to forward to the Concentrator **312** via conventional network **311** for further processing and publication. In addition to carrying out data management policies, Computing Platform **308** may analyze the data archive in order to determine whether beneficial changes in channel management policy ought to be made. When such changes are identified, Computing Platform **308** may forward recommendations to the Concentrator to ensure that the impact of all contemplated policy changes is understood at every substation that may be affected before the policy is put into effect. When changes in the channel management policy go into effect, the fact and nature of the policy change may be added to the data archive as annotations. Such annotations aid in subsequent data analysis because they may explain discontinuities in the archived metrics such as detected signal amplitude or frequency band. Further, the annotations can be used to determine how effective policy changes have been in improving quality-based metrics such as bit error rate, message detection rate, and the like.

[0085] Considering now the devices at the edge of the network, Figure 6 shows an elevation of a single-phase form 2s GLA Smart Meter **601** which is operable to act as a Remote Hub in a Grid Location Aware (GLA) network. Form 2s is standard format for a residential single-phase meter in the United States. Other embodiments of the Remote Hub device may integrate with three-phase meters, or may not be associated with a meter at all, but plug into a 120V or 240V or other voltage outlet located at a building, especially in a microgrid setting where the edge of the network is defined with a higher resolution than in a typical utility service area. In still other embodiments, the Remote Hub device may be integrated with other devices and instruments on the distribution grid, such as voltage regulators, capacitor banks, step-down transformers, and the like.

[0086] A typical Smart Meter has a layered design of circuit boards conformal to the meter housing, such as a cylindrical transparent dome constructed of glass or Lucite. In the depicted embodiment, the dome may be approximately 1.5" taller than usual to accommodate the height of the Edge Transmitter module **604**. However, the height of the meter housing varies from one manufacturer and model to another and does not place the meter in violation of the form standard. Nearest the top of the housing is the Calculation and Display Board **602**, which is part of every Smart Meter. Typically, the display features on the Calculation and Display board such as indicator lights and a digital readout may be accessed by other components in the housing via interface **606**. The Communications module **603** contains the AMI transceiver circuitry and intelligence. In some Smart Meters the components of the Communications module **603** are contained on the Calculation and Display Board, but other smart meters can accommodate multiple types of Communications Modules by placing the communication components on a separate board as shown. Both configurations are common. If on a separate board, Communications Module **603** communicates with logic on the Calculation and Display board **602** via an interface and cable **606**. Other types of component-to-component interfaces are possible internal to the Smart Meter. The Edge Transmitter module **604** contains the long-range GLA edge transmitter and also the short-range PLC transceiver for the TAN communications. Module **604** also has a CPU/microcontroller with nonvolatile memory that hosts and executes the stored programs of the Remote Hub

control logic, controlling the Edge Transmitter and the PLC transceiver, and the interface **606** to the other logic boards **602** and **603**.

[0087] Figures 7a, 7b, and 7c show top, bottom, and schematic views of one embodiment of the Edge Transmitter module. Referring primarily to Figure 7c, the Edge Transmitter Module communicates with other components of the meter **701** via UART (Universal Asynchronous Receiver/Transmitter) **712**. Power is delivered to power supply **703** from AC mains **702**. Power supply **703** provides appropriate low voltage DC power to computing unit **717**, Amplifier Circuit **705**, and Class D Amplifier **707**. Computing unit **717** is a microcontroller processing unit with on-board volatile and non-volatile memory, and is used for both short range and long range digital signal processing and protocols **704**, scaling and filtering **713**, and to drive amplifiers **707** and **705**. PLC Amplifier Circuit **705** and PLC Coupling Circuit **706** are adjusted so that the medium-frequency signal emitted does not propagate outside the TAN, relying on the service transformer's filtering properties at this frequency, as described in more detail below. The coupling circuit **708** for long range transmissions is much more powerful, requiring a special isolation circuit **711** that prevents coupling circuit **708** from resonating with other long-range transmissions and grid noise at the frequencies to which it is sensitive, when the Remote Hub is not transmitting on a long-range channel. The main components of the coupling circuit **708** are transformer **709** and capacitor **710**. According to some embodiments, the Edge Transmitter module may contain a GPS receiver, such as GPS receiver **718** shown in Figure 7c. Alternately, the Edge Transmitter module may have access to the GPS signal from a GPS receiver located on another component of the Remote Hub, such as the Calculation and Display board, or on a mobile computing device used by an Installer. The Remote Hub may use the GPS signal to record its geospatial coordinates, and/or for synchronizing transmissions on a slotted channel so as not to collide with transmissions from other Remote Hubs. Additionally, a Remote Hub may be programmed to report its geospatial coordinates, or a function of its geospatial coordinates, such as the area of the TAN, on an Edge-to-Substation channel or to permit them to be read by a mobile computing device. Remotes whose geospatial coordinates were recorded in a geospatial database (e.g. a PNM) at the time of installation may also receive their coordinates in a data block.

[0088] Figures 7a and 7b illustrate how the components of Figure 7a may be arranged to conform to the shape of a form 2S electrical meter. As is apparent from Figure 7b, transformer **709** and capacitor **710** may be large components. Shown here removed from the assembly to reveal the transformer, flux shield **715** normally covers transformer **709** to prevent magnetic flux from the large transformer from interfering with the metrology unit below it. Referring to Figure 7a, the components of Power Supply **703** occupy a region of the Remote Hub module at the upper left, and the components of Class D amplifier **707** are shown at the upper right. In this embodiment, a single microcontroller **717** contains circuitry, a processor, and nonvolatile memory for firmware protocol stacks and network management logic **704**, **705**, and **713** (from Figure 7c). The large components at the lower right are PLC coupling circuit **706**, isolation circuit **711**, and an amplifier capacitor **716** for the long-range edge transmitter. Connector **712** is the connector for interfacing with the other logic boards in the meter housing.

[0089] Figure 7d details the isolation circuit **711**. This circuit provides the ability to disconnect the coupling circuit of the Edge-to-Substation transmitter from the power line. It is desirable that the Remote Hub be isolated from the power line except when it is transmitting. This is because the coupling circuit represents a substantial load when connected to the power line. The isolation circuit consists of a Triac **718** and a relay **719**. The purpose of the Triac is to allow connection to the power line at the time of a zero crossing of the line voltage. This prevents the generation of large transients that can damage components in the Remote Hub. Additionally, making the connection with the Triac removes the concern of arcing on the relay contacts that reduces the life of the relay. The sequence of events when a Remote Hub interacts with the power line is:

- a. The Triac is closed at a zero crossing of the power line;
- b. The relay is closed;
- c. The desired action (generally a transmission) is performed;
- d. The relay is opened;
- e. The Triac is opened at a zero crossing of the power line.

[0090] Figure 8 is an elevation of a standard Smart Meter **801** with a standard Calculation and Display Board **802** and a standard metrology unit **804**. Meter **801** becomes a Subordinate Remote GLA Smart Meter with the addition of PLC

communication capabilities on the Communications Module **803**. The PLC communications components are small enough that they can share space on Calculation and Display Board **802** with the AMI communications circuitry if preferred. All three modules (if three are present) communicate via an interface, here, a serial interface or UART **805**. Other embodiments may employ a different, functionally equivalent internal inter-component interface.

[0091] Figure 9 is a schematic diagram of the PLC communications components on the Communications Module of one embodiment of a Subordinate Remote. By analogy with similar components on the Edge Transmitter Model in Figure 7c, the communications module of a Subordinate Remote comprises the basic meter **901**, power input from AC mains **902**, interface to the basic meter **907**, power supply **903**, microcontroller unit **904**, amplifier circuit **905**, and PLC coupling circuit **906**. The transmitter may be tuned to ensure that the signals are sufficiently attenuated by the service transformer not to be received above the service transformer or below adjacent service transformers. A Subordinate Remote may also contain a GPS receiver. Alternately, a Subordinate Remote may be programmed with its geospatial coordinates by means of a GPS receiver in a mobile device used by an Installer. A Subordinate Remote that knows its geospatial coordinates may report them over a Transformer Area Network to a Remote Hub. This allows the Remote Hub to compute extended geospatial information such as the area and extent of the TAN. The Remote Hub may report this extended geospatial information on an Edge-to-Substation channel, or the information may be read from the Remote Hub by a mobile device used by an Installer or other field engineer.

[0092] Figure 10a illustrates both the characteristics of the low end of the frequency spectrum on a feeder-phase line of a typical distribution grid and the method employed by a Remote Hub's Edge Transmitter to inject current modulated signals onto the grid in such a way to allow all the TANs served by one substation transformer to be able to transmit scheduled Grid Location Awareness reports at least twice in each 24 hour period, and to additionally transmit alerts as necessary, without creating any of the difficulties described herein above which have been observed with prior art methods of on-grid transmissions. Important characteristics of the spectrum are the 50Hz or 60Hz power fundamental **1001**, its harmonics **1002**, and the noise floor **1003**. It should be

noted that from time to time a spike of impulse noise may exceed the usual noise floor. The defined channel or channels for modulated signals transmitted by an Edge Transmitter occupy a broad candidate spectrum lying between the 50 or 60 Hz power fundamental and the low-pass threshold of the service transmitters on the host power grid. The candidate spectrum for a particular substation is determined by measurement and set by policy and subject to regulatory constraints. Measurements determine which band or bands of the candidate spectrum are reliably received at each substation transformer. If a usable band is wider than the bandwidth needed for a reliable transmission, then the channel band may be defined to be variable. In such cases, the Remote Hub conducts measurements, described herein below, prior to transmitting to determine at the present conditions which part of the wider channel is currently most favorable for transmitting. Conversely, at the Substation Receiver, the preamble detector samples the entire wide usable band, determining the actual band used by the transmitter based on where the preamble was detected.

[0093] Figure 10a shows three frequency bands **1004**, **1005**, and **1006** that have been defined as channels for the long range Edge Transmitters. The number of bands used as channels is not limited to three, nor are three channels always required. Transmissions on each channel are spread across a defined frequency band as shown using a broadband modulation technique, such as the techniques identified herein above. Additionally, transmission bursts may be constrained to occur in time slots such as **1007**. Details of the slotting protocols are explained herein below.

[0094] Still referring to figure 10a, a typical frequency based channel of the present invention may span a wide enough area of the spectrum that several harmonics of the power fundamental occur within the channel. Because it is important to keep the amplitude of the injected, modulated current as near as possible to the noise floor and to minimize the amount of power used to transmit, in some embodiments of the present invention no signal is added to the spectrum at harmonics of the power fundamental. A shaping filter may be beneficially applied by the Edge Transmitter to avoid injecting current over the harmonics. This technique is also beneficial at the Substation Receiver, which may apply comb filtering so that the preamble detector and demodulator are not

required to process the signal on the harmonics. This saves valuable processor capacity in the compute-intensive demodulation process.

[0095] When the modulation technique used is frequency spreading, each frequency band (such as **1004**, **1005**, and **1006**), which is used as a channel, is assigned at least one patterned code, or *chip*. The rate of frequency variation of the chip is much higher than the rate of variation of the data signal. The actual frequency-spread transmission injected as current on the channel is the exclusive or (XOR) of the channel's chip and the data signal. Adjacent and nearby channels are assigned mathematically orthogonal chips. The amplitude of the frequency spread current signals is as close as possible to the noise floor of the power line. This is beneficial in eliminating the problems associated with prior PLC methods. For example, if a transmission on one channel is "folded over" into another channel due to crosstalk, using the different encoding chips causes the receiver to interpret the "stray" signal as noise, allowing the receiver to still extract the correct signal. Additionally, and regardless of the modulation technique, any harmonics from one channel which extend into adjacent channels will also be interpreted as noise. The result of this combination of channel access restrictions and modulation techniques is one or more low-frequency, high-quality current-modulated channels that can bear (in individual bursts) a raw data rate of 120 bits per second or more, or, by example, 80 bps after forward error correction, using interleaving techniques to distribute data bits and FEC bits to minimize the probability of loss of related bits due to impulse noise. Time-duration testing on a radial distribution grid, transmitting an average distance of 3.5 miles line of sight from the substation, yielded a frame error rate of $1.6e-6$ using an FEC rate of $2/3$ with frequency spread modulation. It is recognized that the method and apparatus described may additionally be operated at lower data rates than cited.

[0096] Figure 10b illustrates one method of organizing a group of three reliable channels to support grid mapping. All three channels are organized into time slots **1007**, within which Remote Hubs are provisioned to transmit around a 5-second burst **1008** with around 1 second of silence prior to the burst and around 1 second of silence after the burst. This yields an inter-burst interval **1009** averaging 2 seconds long. Different time intervals may also be used. The reason for the long inter-burst interval in the illustrated embodiment is that the mechanism for synchronizing the transmitter clocks may be an

AMI network, and the AMI synchronization mechanism, being typically based on a mesh or cellular wireless architecture, is no more precise than plus or minus one second. Collisions (overlapping transmissions) on the same channel must be prevented because they will destructively interfere with one another if they were modulated using phase-shift keying or the same chip. In one embodiment of the invention, each data-bearing frequency-spread channel is assigned a plurality of chips instead of one. For example, if the number of chips per channel is two, then transmissions on even-numbered time slots use one chip, and transmissions on odd-numbered time slots use another, mathematically orthogonal chip. Using multiple chips may allow data-bearing capacity of the channel to be increased by reducing the inter-burst interval, as overlaps of adjacent transmissions may still be decoded. The ordinal number of a time slot is determined with respect to a Master Frame Origin, which may be defined as beginning at midnight local time of each day, or may be established by a variety of methods as described in US Patent Application No. 13/566,481 already referenced and incorporated herein.

[0097] Two of the channels **1011** in Figure 10b have a scheduled organization. This means that each Remote Hub is assigned specific time slots in which it may transmit on the channel. A third channel **1012** is still organized into slots, but any Remote Hub with an exceptional condition to report may attempt to transmit in any time slot, provided that it has not alerted recently. Specifically, channel **1012** is organized by the method known as *slotted aloha*. Alerts, when received at the substation, are typically acknowledged via a Substation-to-Edge channel such as a wireless AMI network. If available, other methods for acknowledging alerts may be employed. If no mechanism for acknowledging alerts is available, then each alert can simply be transmitted multiple times, with a randomly selected number of slots having elapsed between the transmissions. This, however, reduces the data bearing capacity of the alerting channel **1012**, because in standard slotted aloha, alerts are retransmitted only if they are not acknowledged. The rate of message failure will be the frame error rate of the channel (already disclosed to be very low) plus the rate of collisions. The rate of collisions in turn depends on the *offered load*, which is based on the probability that more than one Remote Hub will attempt to transmit in a given slot. The optimum number of unacknowledged

retransmissions to maximize message success rate is likely to be a small number such as two or three, because with higher transmission rates channel saturation may occur.

[0098] The organization of an unscheduled channel may also use an un-slotted protocol similar to pure aloha, wherein the channel is not divided into time slots, but wherein a transmitter may attempt to transmit at any time, given that it has not already transmitted within a predefined recent interval. In this organization, alerts may preferably be retransmitted only if not acknowledged within a predetermined period of elapsed time, or they may routinely be transmitted a multiplicity of times if acknowledging alerts is impossible or undesirable.

[0099] The number and organization of channels described is by example only. On some substations, only one reliable channel may be available. When only one channel is used, either because of conditions or by design, a plurality of time slots may be reserved for alerting, while other time slots are scheduled. On some substations, a plurality of reliable channels will be identifiable. The number of scheduled channels needed depends on the number of Remote Hubs and the number of scheduled messages each Remote Hub must send in a 24-hour period. In one embodiment, two channels are sufficient to permit 12,000 hubs to transmit twice daily. If (as is usual) the substation transformer supplies many fewer than 12,000 hubs, fewer channels than are available are required for scheduled messages, alert thresholds may be lowered, and more than one channel may be dedicated to alerts to accommodate the higher offered load. Figure 10b shows four alerts transmitted in the time interval shown. Two of the alerts **1010** have a high probability of being detected at the substation receiver. The alerts **1013** have collided in Figure 10b and will not be received correctly. Figure 10e, conversely, illustrates how the use of chips selected by means of the modulus of the time slot prevents some collisions. Here, because of poor synchronization of clocks, message **1014**, transmitted in an even-numbered slot of random slotted channel **1012**, has overlapped with message **1015**, transmitted on the same channel in the subsequent odd-numbered slot. Both messages are decipherable at the substation because they were encoded using orthogonal chips. In yet another embodiment of the invention, an unscheduled, unslotted channel might use frequency spread modulation and be assigned a plurality of orthogonal chips. A transmitter offering a message would select a chip from the plurality of chips at random,

thereby reducing the probability that the message would collide with another transmission on the same channel at an overlapping time.

[00100] Figure 10c illustrates the detailed organization of a typical single transmission burst, whether it occurs on a scheduled channel or a slotted alerting channel according to one embodiment. Within Time Slot **1007** and burst **1008**, the message is comprised of preamble **1014**, interleaved data bits **1015** and FEC bits **1016**. The preamble is the same for all messages on the channel. The FEC rate is not drawn to scale, and may be varied as needed from substation to substation based on the quality of the available channels. In some grid locations and/or with some modulations, FEC may not be required. Figure 10c without further elaboration may appear to imply that the bandwidth is the same for all transmissions in the same channel, and that the pattern used for preamble detection is also suitable for use by a Substation Receiver when sampling and comparing signals on several inputs representing different phases of different feeders to infer the line on which the signal was actually transmitted. Some embodiments of the invention may require greater bandwidth for preamble detection than the data-bearing segment of a message requires. Additionally, in some embodiments, the grid location of a Remote Hub transmitter may be better inferred from a special transmission, called a probe transmission, again measured at the substation on all phases of all feeders monitored by a Substation Receiver. The probe transmission may consist of known modulated signal, or it may consist of pure tones. The pure tones may be transmitted as a sequence of single tones, or one or more groups of pure tones may be transmitted simultaneously. The frequency range of the probe transmission may be different from that of the other message sections. Figure 10d illustrates this bandwidth variation, showing one bandwidth for the preamble **1018**, another bandwidth for data-bearing message **1019**, and a third bandwidth for GLA trailer **1020**. GLA trailer **1020** is not present in all embodiments of the invention, because the probe transmission may be present within preamble **1018**. In another embodiment, the probe transmission may precede the preamble rather than following the message. Generally, the segments of a message may be transmitted in any order as long as the order is known by the receiver.

[00101] Figure 11 shows a Remote Hub **1101** and a Subordinate Remote **1102** illustrating the local communication paths within a TAN according to one aspect of the

invention. This figure shows an embodiment where the Communications Module is not separate from the Calculation and Display Module. The Remote Hub **1101** may poll each known Subordinate Remote, via a PLC protocol such as PRIME or G3 using request path **1103**. (To allow for the use of different PLC protocols, the specific language of these standards is not used herein. By way of example, if the PLC protocol in an embodiment of the invention were PRIME, then the Remote Hub would be a PRIME base node and all other nodes would be service nodes.) A polled Subordinate Remote **1102** retrieves the requested data from the Smart Meter and formats it into a response which is transmitted as a response **1104**. The Remote Hub's Edge Transmitter Module communicates with the Communications Module and Calculation and Display Board components via UART **1107**, using a simple request/response protocol **1105** that may vary from one Smart Meter vendor to another. Data path **1106** illustrates that both the Remote Hub and the Subordinate Remote are members of the AMI and will be transmitting meter data to the AMI head end in addition to performing TAN-related activities. The Remote Hub, in its role as TAN manager, may make use of the AMI or other alternative, integrated channels in ways that Subordinate Remotes may not. Only the Remote Hub is capable of sending out messages on the Edge-to-Substation channel. A Remote Hub may also send messages on alternative, integrated outgoing channels such as the AMI. The Remote Hub may additionally receive data blocks from a Substation-to-Edge channel, whether the Substation-to-Edge channel is provided by the AMI or other means. Such data blocks may contain, but are not limited to, alert acknowledgements, firmware update broadcasts, and policy changes. Meter clock synchronization messages are part of the native AMI protocol, but the Remote Hub may obtain the synchronized clock time from the Calculation and Display module when an AMI is present.

[00102] Remote Hub **1101** has the capability to function in multiple operating modes. The Remote Hub may function as a Subordinate Remote. The Remote Hub may also function as a hybrid of Remote Hub and Subordinate Remote, called a Proxy Hub. When a Remote Hub **1101** is first installed, it monitors the PLC frequencies on the TAN for a period of time sufficient to determine whether another Remote Hub is already present. The wait time consists of a fixed period of time plus an additional period of time computed by a randomization function when the device is powered on. The fixed period

of time is sufficient to ensure that a Remote Hub operating in the master mode would have executed its discovery algorithm, which would be detected by the newly installed Remote Hub if another Remote Hub is operating within range. Typically, “within range” means powered by the same service transformer, but exceptions occur. The means of handling the exceptions are described herein below.

[00103] If a first Remote Hub is already present, Remote Hub **1101** indicates by means of a light or digital display on the face of the Smart Meter that another Remote Hub is present. At this point, an installer may elect to leave redundant Remote Hub **1101** in place, or replace it with a Subordinate Remote unit. If left in place as a “spare,” Remote Hub **1101** continues to function as a Subordinate Remote, and the first Remote Hub continues to act as the Remote Hub and master node in the TAN. If no other Remote Hub is present, Remote Hub **1101** begins to operate as a master PLC node on the TAN, discovering and storing a list of any Subordinate Remotes **1102** in the same TAN. A Remote Hub may also enter a third mode, Proxy Hub, as described below. As soon as it takes on the master or hub role, Remote Hub **1101** obtains the network system time if available, for example by querying the AMI logic in the Smart Meter, and formats, encodes, and transmits a provisioning request on an Edge-to-Substation channel reserved for provisioning requests and alerts. When a Substation Receiver detects the provisioning request, it may cause a provisioning response to be sent, either via the AMI, or via an available on-grid or wireless Substation-to-Edge channel. Provisioning data may also be supplied to a Remote Hub by means of a handheld device or drive-by transmitter employed by the installer. The handheld device uses a personal-area wired or wireless protocol, such as Bluetooth, infrared, USB, or RS232 to program the Remote Hub. In embodiments of the invention where the Substation-to-Edge channel is absent or very constrained, the Remote Hub may be provisioned via handheld without knowledge of the inferred grid location of the Remote Hub. The same short-range protocol, in a handheld or drive-by device, may be used to distribute firmware or policy updates to Remote Hubs that lack a permanent Substation-to-Edge channel. It is sometimes desirable to activate a policy or program change simultaneously on a collection of Remote Hubs. If the Remote Hubs must be updated by means of a personal-area protocol, the programming device converts the desired future activation time to a relative wait time as each Remote Hub is

programmed, so that even though the Remote Hubs were programmed at different times, they will activate the updated programming at approximately the same future time.

Remote Hubs may be manufactured with a default policy, or pre-loaded with a default policy after manufacturing but before installation, so that if no policy is provided at or subsequent to installation, the Remote Hub still has a rule for operating.

[00104] The provisioning data provides the Remote Hub with the information it needs to manage the TAN, including the location of, and organization of, other channels on the Edge-to-Substation network, and the ordinal or sequence number of slots on scheduled channels on which this Remote Hub has permission to transmit. When the Remote Hub discovers Subordinate Remotes, it transmits pairing messages on the Edge-to-Substation channel to inform the Computing Platform that it is in communication with the newly discovered Subordinate Remotes. Pairing messages may be transmitted on an alerting channel or on a scheduled channel depending on a policy established by the network. When a Remote Hub acting in the master role has discovered another Remote Hub on the same transformer and phase operating in the subordinate role, the resulting pairing message indicates this. Including the presence of “spare” Remote Hubs in the Grid Map may provide a cost savings and more rapid recovery, in that if the master Remote Hub should fail, the TAN may be reconstructed by allowing the spare Remote Hub to assume the master role. The master Remote Hub may cache its policy information on a spare Remote Hub, if present, in order to allow the failover to occur even without re-provisioning the TAN.

[00105] Hereinafter are disclosed methods for properly partitioning Remote devices into TAN groupings. PLC transmission power is controlled in order to keep the signal that gets through the Service Transformer low enough to avoid interference with other TANs. Specifically, unless special accommodations in configuration are made as described herein below, a Remote Hub must poll and collect data from only Subordinate Remotes on the same phase of the same service transformer as the Remote Hub. However, at certain sites on some grids, it may happen that at PLC standard power and frequencies, the PLC transceivers in the Remotes may be able to discover Subordinate Remotes and Remote Hubs on other phases of the same service transformer, or even on adjacent or nearby service transformers. In this aspect of the invention, the detectable

remotes may be partitioned wherever possible so that each TAN comprises exactly one master Remote Hub and all Subordinate Remotes, or Remote Hubs acting as Subordinate Remotes, on the same phase of the same service transformer, and no Remotes of any type which are on a different phase or a different service transformer.

[00106] In one embodiment of the invention, a Remote Hub's PLC protocol stack executes its discovery process, which involves transmitting a beacon tone or message that causes other nodes in the vicinity to respond. The first time this is executed, a standard initial power level is used. The TAN management layer of the Remote Hub, operating above the PLC protocol stack, obtains the list of discovered Remotes of any type. The Edge Transmitter of the Remote Hub is then employed to send a pilot signal at sufficiently low amplitude and high frequency that the pilot signal will not be detectable on the high-voltage side of the service transformer. (This pilot tone is not the same as a PLC discovery beacon.) The pilot signal begins on a zero crossing of the power fundamental of the phase on which the transmitter resides. Other Remotes (of any type) that detect the pilot signal test to determine if the received signal began on the zero crossing of the phase on which the receiving Remote resides. If so, the receiving Remote sends a positive response on the PLC channel and records the identity of the Hub Remote that sent the pilot tone. Another Remote Hub on the same phase as the pilot transmitter enters Subordinate Remote mode and will be considered a spare. Subordinate Remotes on other phases do not respond to the pilot tone. A Remote Hub that detects the pilot tone but is on a different phase sends a negative response. The transmitting Remote Hub uses the responses to update its inventory of TAN devices discovered automatically by the PLC discovery process, recording the list of Subordinate Remotes and spares on its home phase, and the list of Remote Hubs on other phases of the same service transformer. Remote Hubs that sent neither a negative nor a positive response are presumed to be on another service transformer. If this case exists, the value of the "initial" power level (amplitude) for the PLC discovery beacon is reduced, so that next time the full discovery process is executed, it will be less likely that any Remotes on other service transformers will respond.

[00107] Next, the first Remote Hub that transmitted the pilot tone examines the list of negative responders, that is, of Remote Hubs on a different phase. It selects one such

second Remote Hub and orders it via the PLC protocol to transmit a pilot tone of its own. The first Remote Hub, still the master node of at least all the nodes on the service transformer, collects the resulting positive and negative responses and updates its inventory and partitioning data. At this point, any spare Remote Hubs on the same phase as the second Remote Hub have also entered Subordinate Remote mode, and the first Remote Hub now has a complete partitioning of Remote Hubs according to phases, the Remote Hubs on the third phase, if present, being the intersection of the Remote Hubs sending negative responses to the first Remote Hub with the Remote Hubs sending negative responses to the second Remote Hub.

[00108] If a third phase is present, the first Remote Hub now selects a third Remote Hub from the third phase, and orders it via the PLC protocol to transmit a pilot tone and return the list of negative and positive responses it received. At this point, a positive response will have been received from every Subordinate Remote on the service transformer, the phase and mode of every device on the service transformer is known, and a potential master Remote Hub for each single-phase TAN has been identified. Additionally, any node that responded to the original PLC discovery process from outside the transformer area has been identified.

[00109] Now the first Remote Hub sets its PLC transmission amplitude to a very low level and polls each remote. This first amplitude should be so low that no remotes respond. The first Remote Hub increases its transmission amplitude until, ideally, all remotes on the same phase and no remotes on another phase respond. The first Remote Hub records this low threshold level and then continues to increase the amplitude until a remote on another phase responds. The first Remote Hub records this as its high threshold level.

[00110] Now the first Remote Hub commands the second Remote Hub via PLC to attempt to take on the role of PLC master node for its phase, sending in the command the low and high threshold amplitudes. This is called the partitioning command. The second Remote Hub sets the PLC transmission amplitude to the low threshold amplitude, and restarts its PLC stack as a master node, conducting a PLC discovery process of its own. If the second Remote Hub discovers all the Subordinate Remotes and spares on its own phase and no nodes on any other phase, then has become the master of a single-phase

TAN and the partitioning step has succeeded. Otherwise, it raises its PLC transmission amplitude and repeats the process until the partitioning step succeeds. If the second Remote Hub reaches the High Threshold amplitude without having discovered all the Remotes on its phase, or if at any amplitude a Remote from a different phase is discovered when no lower amplitude discovers all the Remotes on the same phase, then the partitioning command has failed. The second Remote Hub signals the failure of the partitioning command to the first Remote Hub by using its Edge Transmitter to transmit a status beacon detectable by the First Remote Hub, since the first and second Remote Hubs can no longer communicate via PLC.

[00111] If the first Remote Hub detects no failure beacon from the second Remote Hub, and a third phase is present, the first Remote Hub sends a partitioning command to the third Remote Hub, which carries out the partitioning step as described.

[00112] When the first Remote Hub has partitioned the other phases present without having received a failure beacon, then it carries out the partitioning step itself. If the first Remote Hub's partitioning step succeeds, then the service transformer is successfully partitioned into three single-phase TANs, as shown in Figure 12a. In another embodiment of the invention, the second and third Remote Hubs may employ a failure beacon and a success beacon. Use of the success beacon may shorten the time required to complete the partitioning steps.

[00113] Referring now to Figure 12a, which is a simplified schematic drawing of a three-phase service transformer and the meters that it supplies with power. This service area contains three TANs **1204**, **1205**, and **1206**, one for each phase of the service transformer. Each TAN contains a Remote Hub **1202** and zero or more Subordinate Remotes **1203**. Any Subordinate Remote may actually be a spare Remote Hub. Figure 12a illustrates a proper partitioning of the Remotes powered by a three-phase transformer following the discovery and partitioning algorithm described herein above.

[00114] As is clear from the above description of a discovery and partitioning algorithm, it is possible that for some multi-phase transformers there is no set of PLC transmission frequencies that will yield a clean partitioning of the Remotes on the service transformer into single-phase TANs. When the partitioning algorithm fails at any step, the first Remote Hub attempts to form a multi-phase TAN which includes all Remotes on

all phases of the service transformer, but no Remotes which are not on the service transformer. Refer now to Figure 12b, which illustrates a multi-phase TAN. Recall that the first Remote Hub already has an inventory of all Remotes on any phase of the service transformer, and that it further is aware which nodes on each phase are Remote Hubs. Beginning with the previously recorded “initial” PLC transmission amplitude, the first Remote Hub initiates a PLC discovery process. If any Remotes are discovered which are on a different service transformer, the first Remote Hub lowers the PLC transmission amplitude, superseding the old value of the “initial” amplitude, and restarts the discovery process, repeating this until all and only the Remotes known to be on the service transformer are discovered. If a new Remote never before detected is found, the pilot beacon method above is used to determine the phase of the new Remote and whether it is on the same service transformer as the first Remote. If no transmission amplitude can be found that discovers all and only Remotes on the same service transformer as the first Remote Hub, the first Remote Hub transmits a distress alert on an Edge-to-Substation channel and organizes the TAN at the highest amplitude which does not discover any nodes outside the service transformer, even if some nodes on the service transformer are unresponsive.

[00115] For Grid Location Awareness and the energy management applications that depend on the Grid Map to be effective, probe transmissions must originate from each phase of the Service Transformer. To accomplish this, the first Remote Hub, Master **1202** in Figure 12b, sends commands to the second (and third, if present) Remote Hubs **1208**, causing them to operate as Proxy Hubs. A Proxy Hub behaves like a Subordinate Remote on the TAN, except that it is responsive to certain commands from its master Remote Hub that allows the Remote Hub to control the Proxy Hub’s Edge Transmitter. Remote Hub **1202** stores the Edge-to-Substation provisioning policies that would normally be carried out by the Proxy Hubs **1208**. Remote Hub **1202** carries out all the TAN management activities, such as polling the Subordinate Remotes, distributing updates, and computing derived results, for the Remotes, including Proxy Hubs, on all phases present. When it is time for a Proxy Hub to send an Edge-to-Substation transmission, the first Remote Hub **1202** formats the appropriate message and sends it to the Proxy Hub over the TAN. The Proxy Hub then retransmits the message on the Edge-

to-Substation channel. In this way, Edge-to-Substation transmissions are always transmitted on the correct phase, even though the TAN master is on a different phase.

[00116] The partitioning and discovery methods disclosed above are designed to accommodate a standards-based PLC protocol stack such as PRIME. Use of alternative short-range PLC protocol stacks may require minor modifications to the methods. More straightforward methods may also be used in cases where customizations to the lower layers of the protocol stack are allowable.

[00117] In some embodiments of the invention, a Subordinate Remote rather than a Remote Hub may take on the role of the TAN master. In such embodiments, Remote Hubs behave like Proxy Hubs at all times, including during the discovery and partitioning process. Proxy Hubs must include functions to allow the TAN Master to control its Edge Transmitter to send pilot signals during the discovery and partitioning process. This is not fundamentally different from allowing a first Remote Hub to control the Edge Transmitter of a second Remote Hub to send pilot signals. Similarly, any Subordinate Remote must be able to accept a command from its current TAN Master to become a TAN Master itself and attempt to form a TAN within a local PLC frequency range indicated within the TAN. The TAN Master can transmit its own beacon signals, because the local PLC transmitter is used for that purpose. Using this division of function requires modifications to the partitioning and discovery methods so that the TAN Master related functions are assigned to a first Subordinate Remote, and, in cases of multi-phase transformers, a second and possibly a third Subordinate Remote, while leaving the Edge Transmitter related functions with the first, second, and third Remote Hubs. A Remote Hub that is the only Remote in a transformer area will implement its policy regarding Edge-to-Substation transmissions by sampling only its local instruments. It will additionally continue to monitor the PLC frequency for an invitation to join a TAN in case a Subordinate Remote is subsequently installed in the transformer area.

[00118] Another aspect of a Remote Hub's channel management capability is that the Remote-hub may pre-modulate and store certain messages that do not contain variable data and may be sent repeatedly. Examples of pre-recordable messages include messages sent on the provisioning channel, such as the provisioning request and standard alerts on conditions such as sags, over-voltages, and the like. This strategy saves

computing power at the Remote Hub. When policy changes such as changes in chip, channel placement, baud rate, FEC rate, and bandwidth occur, pre-modulated recordings may need to be discarded and re-computed. This may be done during idle periods when the Edge Transmitter's microcontroller CPU is not busy with preparing scheduled messages. To accommodate this, such policy changes may be announced in advance to take effect at a known future time as opposed to becoming effective immediately.

[00119] In some embodiments, a Remote Hub may not be integrated into a GLA Smart Meter, but instead may be associated with another feature of a medium voltage distribution grid, such as a capacitor bank, step-down transformer, voltage regulator, storage battery, local generator, or switch set. The Remote Hub may be integrated with local or remotely controlled SCADA systems associated with the feature. The SCADA system may provide an Edge-to-Substation channel for provisioning Remote Hubs used in this manner, or the Edge-to-Substation channel associated with Remote Hubs in Smart Meters may also be operable to communicate with such feature-based Remote Hubs. Such Remote Hubs may incorporate a version of a Substation Receiver and be operable to send Pairing Messages associating the grid feature with other Remote Hubs electrically and schematically subordinate to the grid feature. A Remote Hub may also be embodied as a stand-alone device plugged into an electrical outlet. A form of Substation Receiver may additionally be associated with such medium voltage grid features, or any intermediate point on the medium voltage distribution grid. Such an intermediate Receiver may collect information regarding which Transformer Area Networks are impacted by an associated medium-voltage grid feature. The combination of such a secondary Receiver and Remote Hub may be employed to control intermediate grid features, such as using a switch or relay to isolate a microgrid or balance the load on a plurality of substations, or to alter the set-point on a voltage regulator.

[00120] In a further aspect of the invention, a Remote Hub may carry out line measurements to determine locally optimum conditions for transmitting. The Remote Hub always has the option to vary the amplitude of the injected signal, and may additionally have the option to vary the frequency band of the preamble and/or data bearing segment of the transmission. Refer now to Figure 13, which shows the same portion of the frequency spectrum as in Figure 10a, but with different annotations. The

bracketed frequency range **1309** spans the entire range of frequencies usable for Edge-to-Substation transmissions. Ranges **1304**, **1305**, and **1306** show how this frequency range **1309** is divided into three frequency sub-ranges, one for each of three channels defined by policy. Within each channel frequency range, the actual data transmissions (**1301**, **1302**, **1303**) may typically use only a portion of the bandwidth assigned to the channel of the transmission. By contrast, the calibration signal described herein below, as well as the locator patterns described above, may span a wider frequency band. Additionally, if the locator pattern and message preamble are different, the message preamble may also span a wider frequency band. The wider frequency band may be the entire range allocated to the channel (such as range **1304**), or even, in some cases, the entire range **1309**.

[00121] To conduct the measurements, the Remote Hub transmits a calibration signal, which may be a sequence or simultaneous combination of pure tones. These tones may be independent of an actual message transmission, or they may be incorporated in the message preamble. Recall that the bandwidth of the preamble may be different than the bandwidth of the data-bearing segment of the transmission. If there is an opportunity to choose the frequency band of the data bearing segment, then the tones must span the entire available range of spectrum. Current is measured in the calibration signal, and is used to calculate a drive voltage that will yield the desired level for the data transmission. When the tones are transmitted, the current generated at the requested voltage is measured. The relationship between the requested voltage and the generated current is calculated at each frequency. The result will be proportional to the line impedance of the grid at the Remote Hub for each frequency. This allows the Remote Hub to estimate both how much drive voltage is required to generate the desired current at each frequency in the available frequency band, and, if there is a choice of frequency ranges to use, to select the frequency range that requires the least voltage to achieve the desired current. The Remote Hub then adjusts its Edge Transmitter configuration to reflect the selected drive voltage and frequency range. The starting points around which these calibrations occur are the base drive voltage and the base frequency range. Referring again to Figure 13, arrow **1307** illustrates how the data transmission frequency range has been adjusted within the channel 1's allowed frequency band, and arrow **1308** show how the amplitude

of the data signal in channel 3 has been increased by increasing the drive voltage. Note that in Figure 13 the offset of the waveform representing the adjusted Channel 1 band after adjustment **1307** is purely for readability of the drawing, and does not represent a vertical shift in amplitude of the signal. Typically, beneficial adjustments will include raising rather than lowering the drive voltage, but the Remote Hub may be configured with maximum and minimum allowed drive voltages so that the specifications of the Edge Transmitter are respected. In some embodiments where a Substation-to-Edge channel is available and has sufficient capacity, the Computing Platform may from time to time send feedback from the Substation Receiver about the messages as received. This feedback is typically based on trend analysis made by analyzing the archived historical metrics described herein above. Trend analysis may be performed via linear or higher-order regressions, by correlating metrics not directly available to the Remote Hubs, such as incidences of impulse noise or variations in the Gaussian noise levels according to the season, the ambient temperature or other weather conditions, or the time of day, against the message success rate for a channel, an individual Remote Hub, or a collection of Remote Hubs based on the topology of the distribution grid. The feedback may cause the Remote Hub to make the one or both of the same types of adjustments that are shown in Figure 13 as frequency band adjustment **1307** and amplitude/signal strength adjustment **1308**. The feedback may also allow the Remote Hub to refine and calibrate its measurement process, such as changing the baselines or defaults that are used as starting points for the self-calibration process described herein above. For example, the archive data may show that a Remote Hub's assumptions about which base drive voltage or target current level results yields the desired signal amplitude at the receiver. If the actual signal amplitude at the receiver is too low for good reception, then feedback might cause a Remote Hub to increase its base drive voltage. Similarly, if the archive indicates that messages transmitted at the higher end of the channel's available frequency band have been received with a lower bit error rate than messages transmitted at the lower end, then a Remote Hub might be instructed to re-center its base frequency band. Outcomes of this feedback may also include changing the slot assignments and/or modulation methods of individual Remote Hubs and/or an entire channel to improve message success rate. This may also include adjusting the frequency band permitted by policy for an entire channel

and all the Remote Hubs which transmit on that channel. Similarly and by example, if historical analysis demonstrates that Channel 1, using QPSK modulation, has a better experience than Channel 3, using frequency-spread modulation, then the Channel 3 policy might be changed for all Remote Hubs to use QPSK as well. Additionally, a Remote Hub may have local hardware and firmware diagnostic processes in addition to the above-described local optimization technique. If one Remote Hub has a historical performance which is significantly different than the performance of other Remote Hubs on the same geospatial or schematic segment of the grid, then that Remote Hub might be instructed to run diagnostics and recalibrate itself. If this process does not bring the Remote Hub's performance into alignment with its neighbors, then a human engineer might be dispatched to the site to investigate, and make line repairs or replace the Remote.

[00122] The foregoing description of the invention has been presented for purposes of illustration and description and is not intended to be exhaustive or to limit the invention to the precise forms disclosed. Obviously many modifications and variations are possible in light of the above teaching. The embodiments were chosen and described in order to best explain the principles of the invention and its practical application to thereby enable others skilled in the art to best utilize the invention in various embodiments and with various modifications as are suited to the particular use contemplated. It is intended that the scope of the invention be defined by the claims appended hereto.

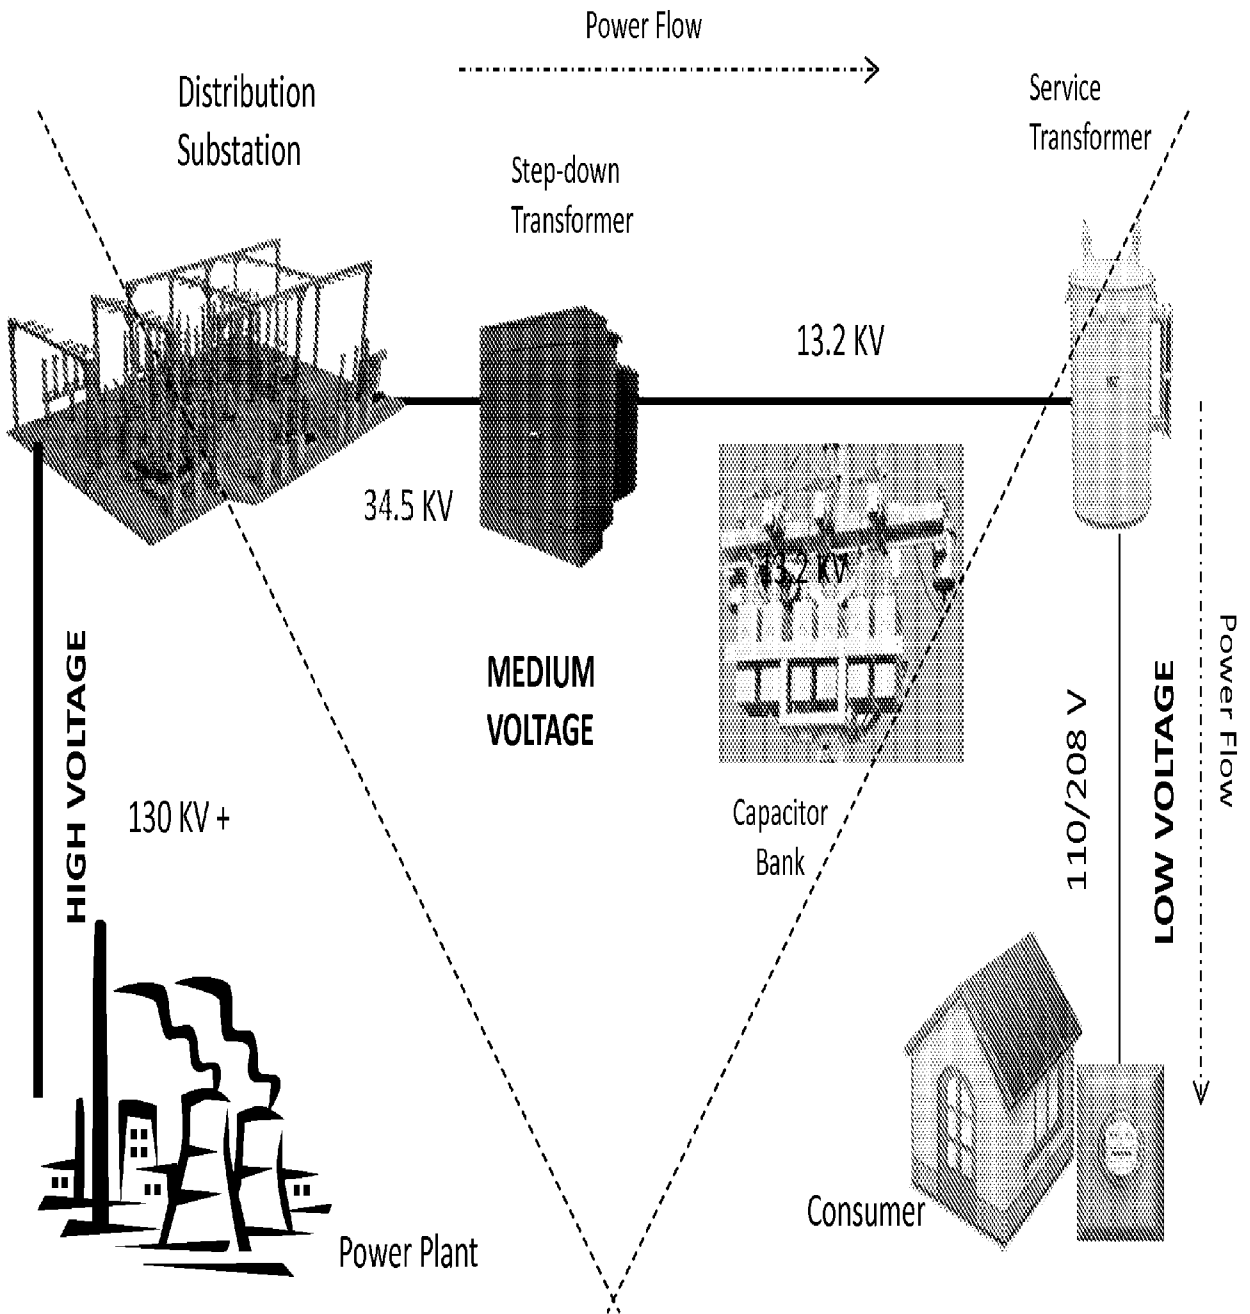
Claims


1. A method for optimizing performance of Edge-to-Substation channels in a communication network utilizing an electrical distribution grid, , the method comprising:
 - a. retrieving archived data about the history of at least one Edge-to-Substation channel;
 - b. analyzing metrics on the archived data relating to the performance of the at least one Edge-to-Substation channel over time; and
 - c. identifying potential improvements to network performance on the channel.
2. The method of claim 1, wherein said metrics include at least one of the ratio of energy per bit to noise power spectral density, message detection rate, bit error rate, impedance variation, crosstalk, Gaussian and impulse noise, symbol constellation separation, inter-symbol interference, receive time drift, or drift in modulation frequencies.
3. The method of claim 1, wherein said analysis includes comparing the historical performance of a plurality of alternative frequency bands.
4. The method of claim 1, wherein the improvements to network performance comprise a policy change on at least one Remote Hub.
5. The method of claim 4, further comprising the Remote Hub activating policy changes upon receipt, at a predetermined future time, or after waiting for a period of time described by the command or directive.
6. The method of claim 1, wherein the improvements to network performance comprise replacing at least one Remote Hub.
7. The method of claim 1, wherein the improvements to network performance comprise recalibrating a transmitter on at least one Remote hub.
8. The method of claim 1, wherein the identified improvements to network performance are encoded into a command or directive executable by a Remote Hub.
9. The method of claim 8, further comprising distributing said command or directive to at least one Remote Hub by a Substation-to-Edge Channel.
10. The method of claim 8, further comprising distributing policy changes to at least one Remote Hub by a local communication device and interface.

11. The method of claim 1, wherein said analysis is executed on a substation Computing Platform.
12. The method of claim 11, further comprising forwarding the outcomes of said analysis over a conventional network to a Concentrator for further analysis.
13. The method of claim 1, wherein said analysis is executed on a Concentrator.
14. The method of claim 1, where the outcome of said analysis is a maintenance work order for at least one Remote Hub.
15. A method for optimizing the drive voltage and selected frequency band for transmitting messages comprising a preamble and a data-bearing segment by a Remote Hub on an Edge-to-Substation channel, the method comprising:
 - a. providing a Remote Hub on an Edge-to-Substation channel within an on-grid communication network;
 - b. transmitting at least one calibration signal at selected frequencies across the available portion of the spectrum;
 - c. measuring the current generated at a drive voltage for each frequency, and computing a relationship of drive voltage to resulting current;
 - d. adjusting at least one parameter of a data-bearing transmission based on the measured and computed relationship of drive voltage to resulting current.
16. The method of claim 15, further comprising adjusting the drive voltage for the data-bearing transmission to produce the desired current.
17. The method of claim 15, further comprising selecting a data-bearing frequency band, within a wider frequency range determined by policy, to achieve desired current levels with a minimum drive voltage.
18. The method of claim 15, further comprising providing feedback from at least one of a Computing Platform and a Concentrator to said Remote Hub.
19. The method of claim 18, wherein the feedback is programmed into said Remote Hub by an installer or field engineer using a device with a personal area network interface.
20. The method of claim 18, wherein the feedback is transmitted to said Remote Hub by means of a Substation-to-Edge Channel.

21. The method of claim 18, wherein said Remote Hub adjusts its base drive voltage or target current level based on said feedback.
22. The method of claim 18, wherein said Remote Hub changes a modulation method, a frequency band of a channel, or its slot assignments based on said feedback.
23. The method of claim 18, wherein said Remote Hub recalibrates its transmitter based on said feedback.
24. The method of claim 15, wherein the at least one calibration signal comprises a sequence of individual tones or a group of tones transmitted simultaneously.
25. The method of claim 15, wherein the at least one calibration signal is incorporated into its own segment of a message, separate from the preamble and data-bearing segment.
26. The method of claim 15, wherein the preamble contains the at least one calibration signals as part of a preamble pattern.
27. The method of claim 15, wherein the calibration signal is transmitted independently of any message.

FIGURE 1.
(background)










FIGURE 2b.
(background)





6/24

FIGURE 3.




FIGURE 4.

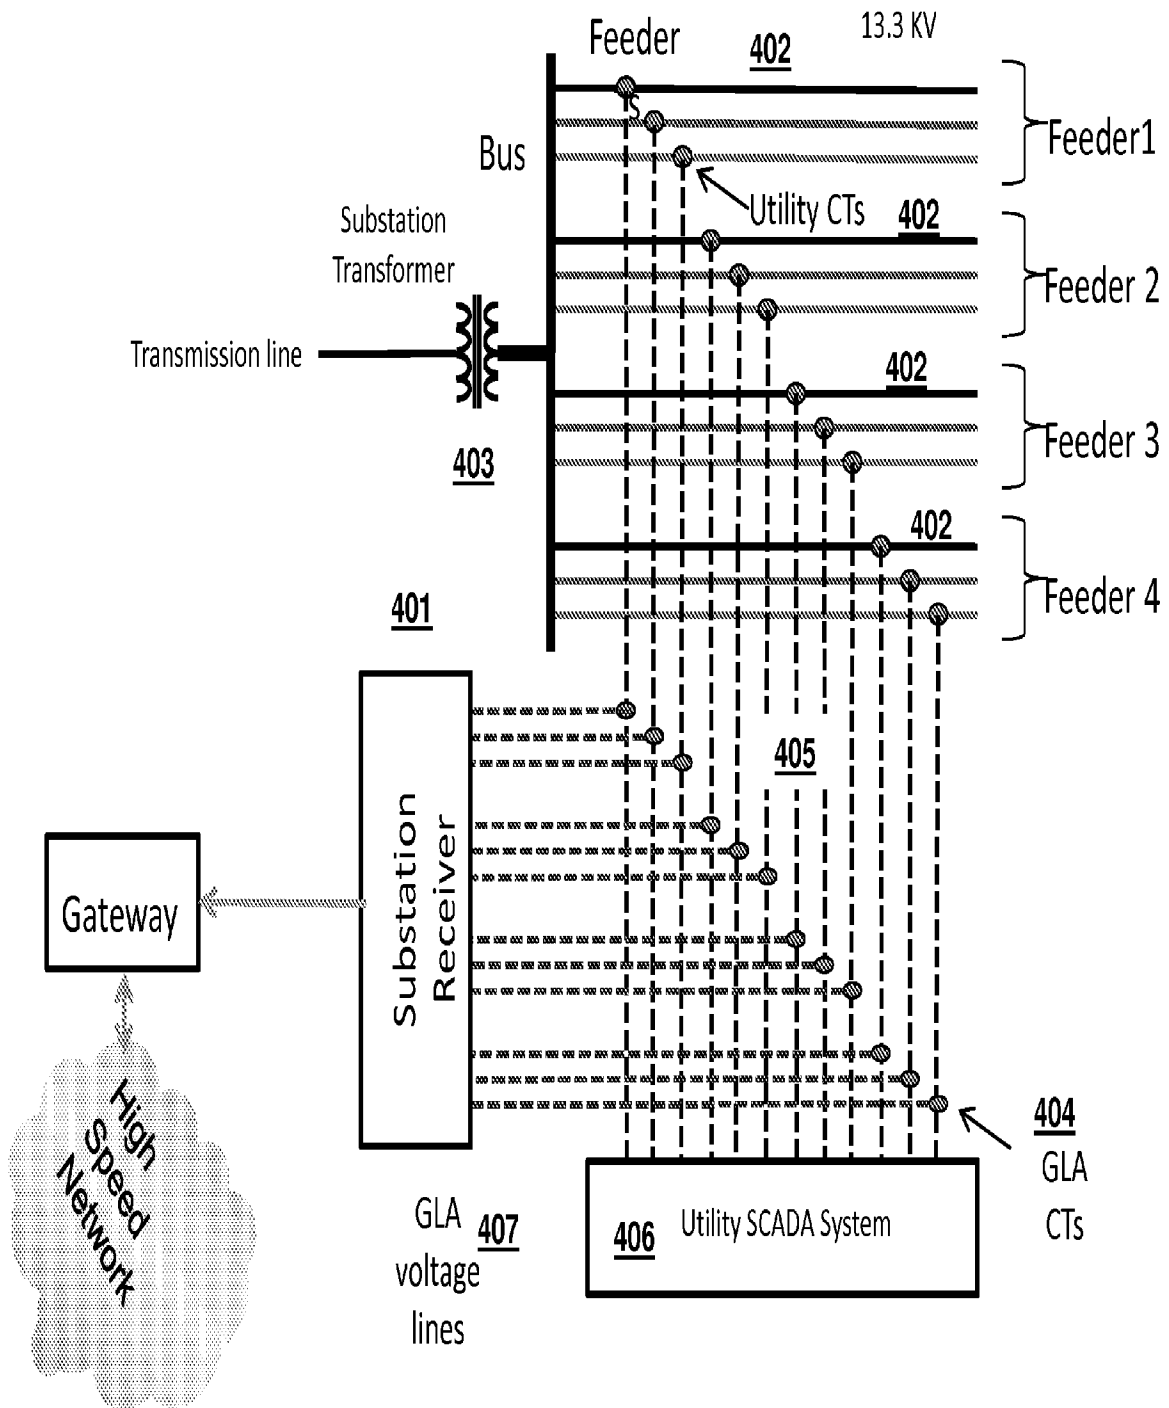
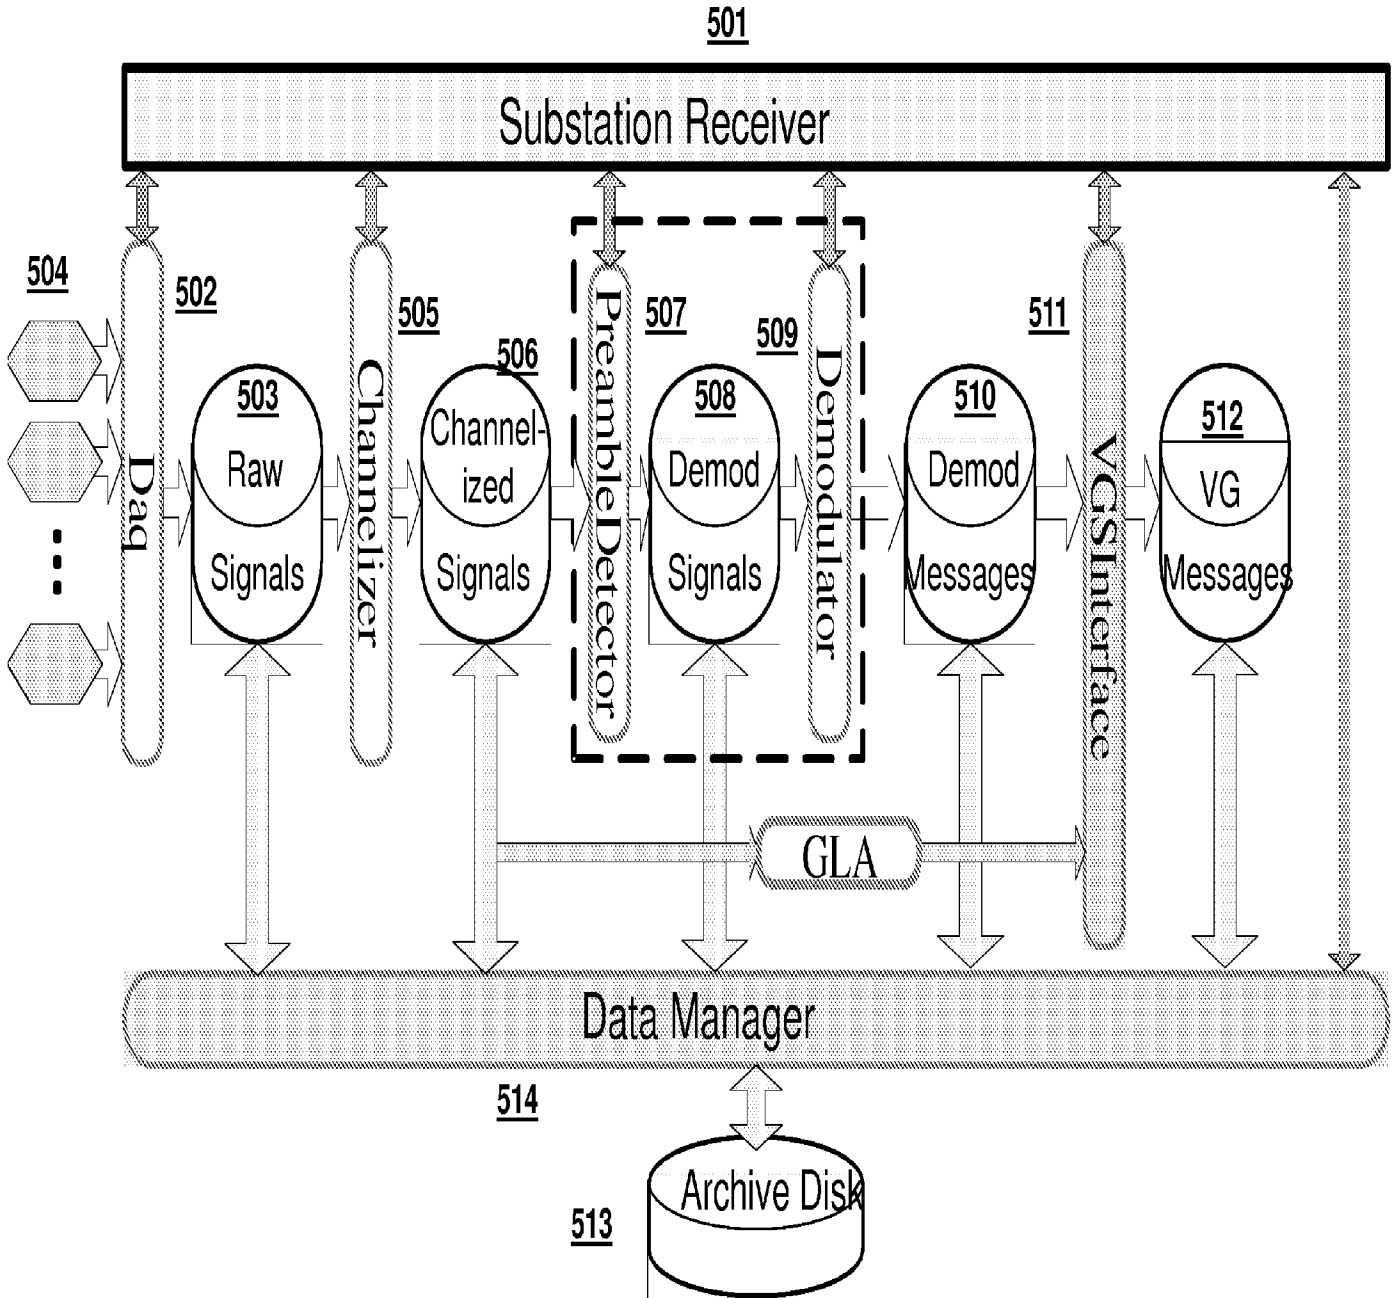



FIGURE 5.




9/24

FIGURE 6.



10/24

FIGURE 7a.



11/24

FIGURE 7b.




FIGURE 7c.

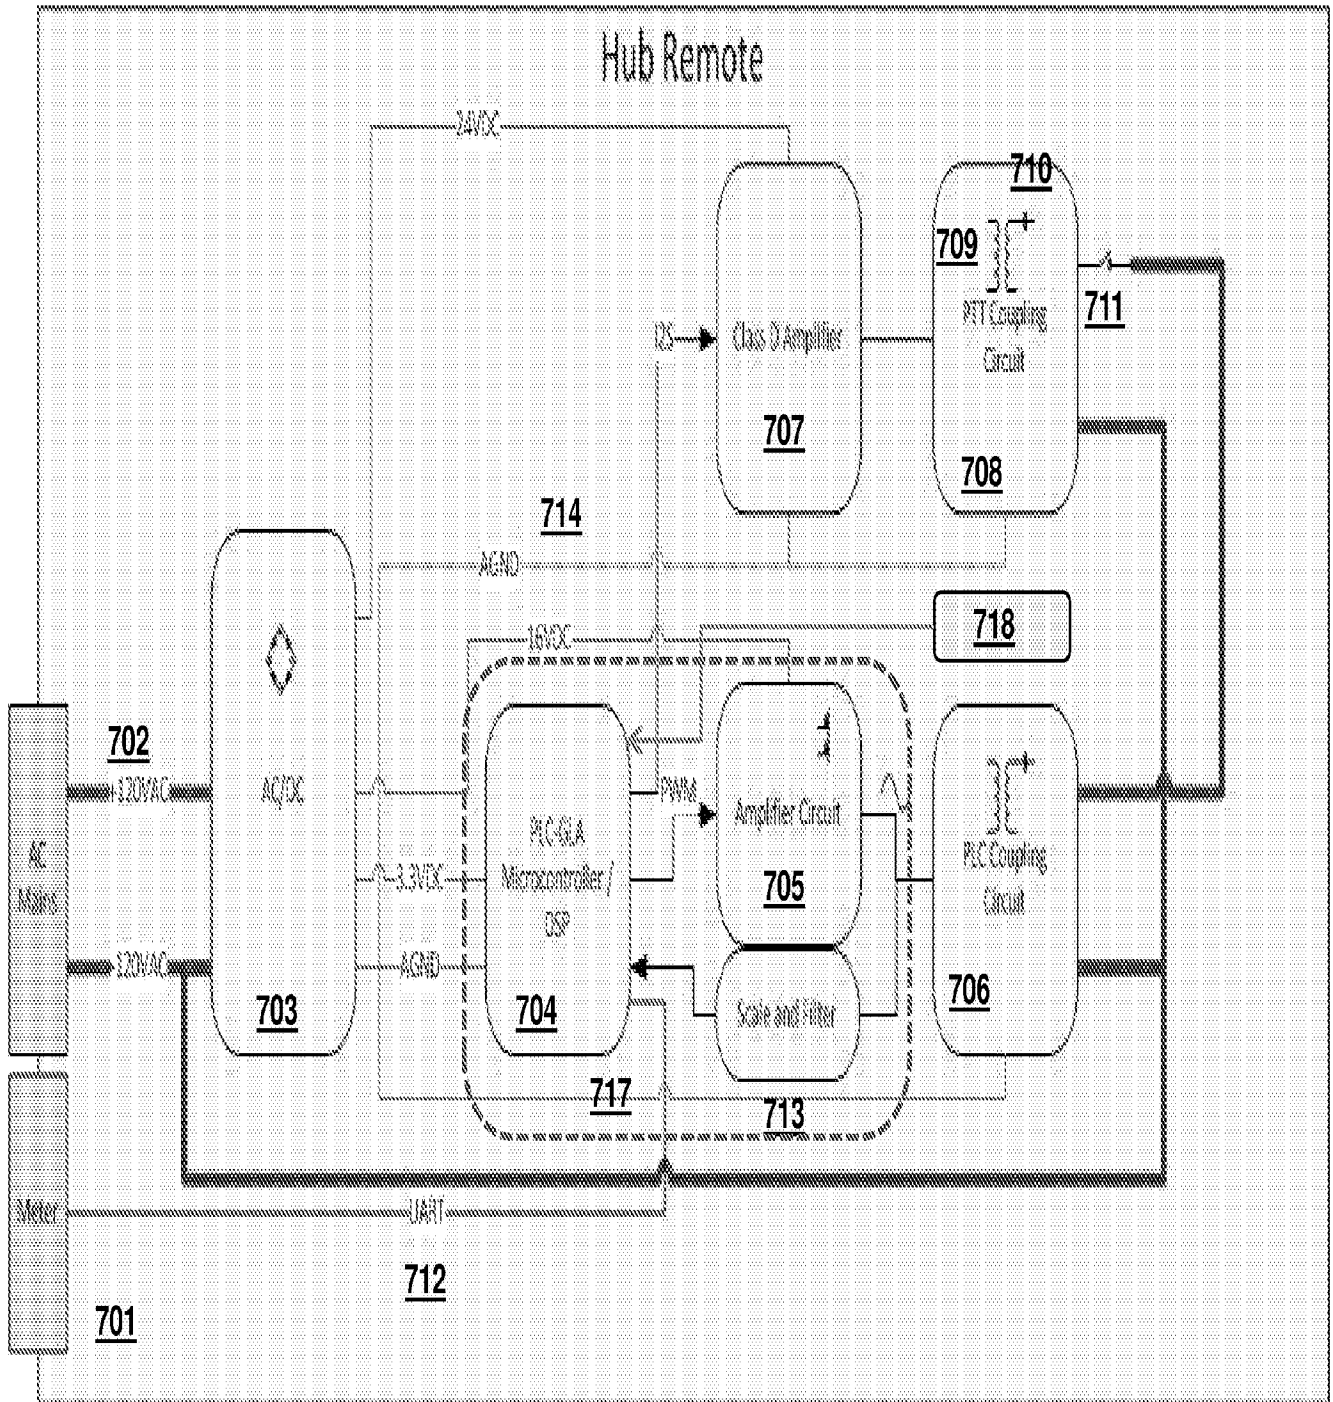


FIGURE 7d.




FIGURE 8.





FIGURE 9.





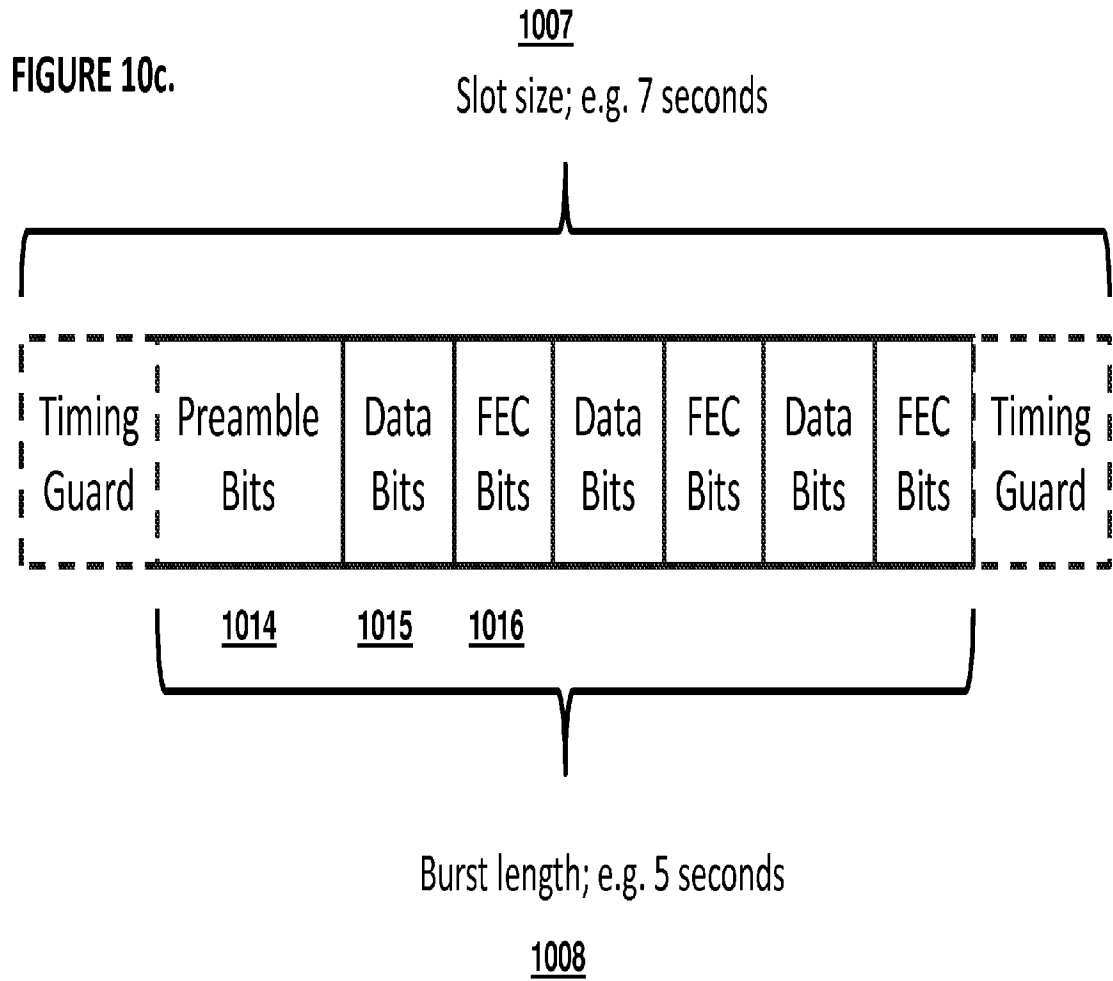


FIGURE 10b.

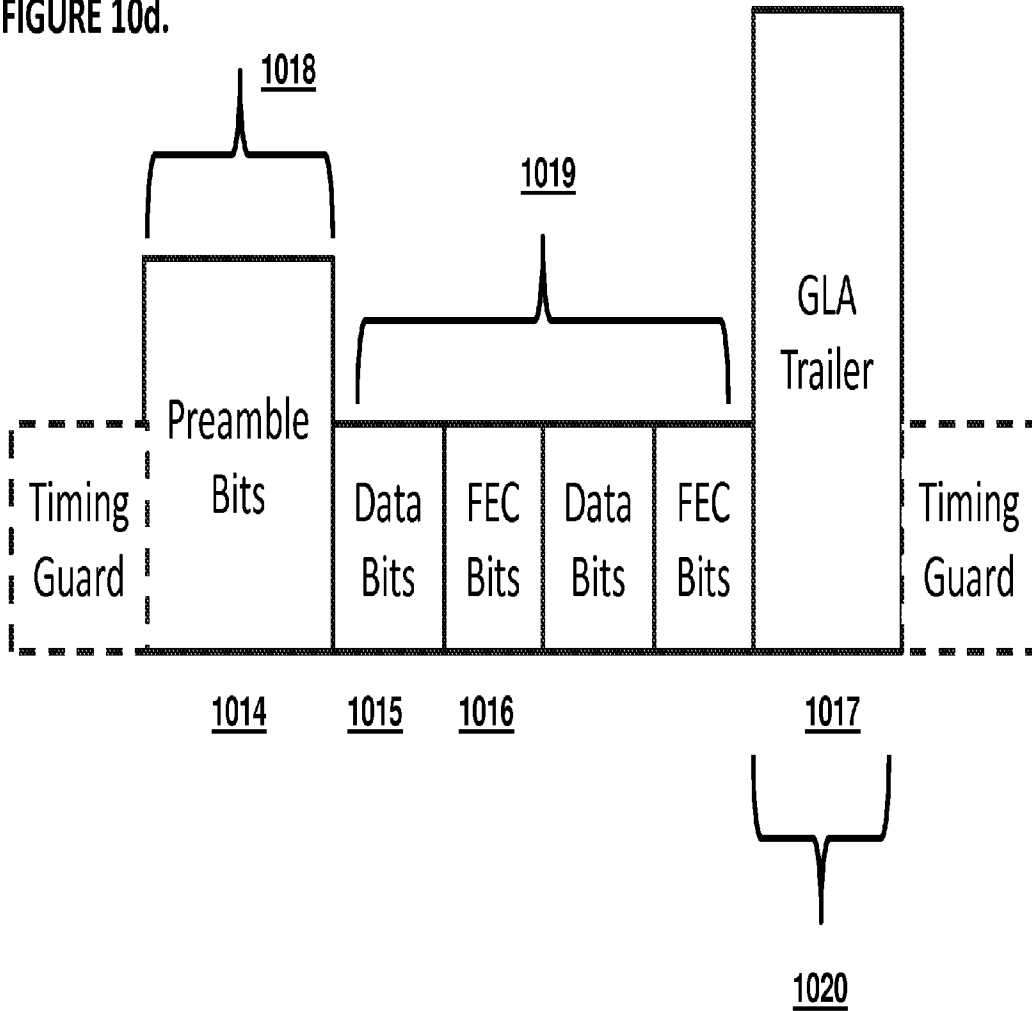


HR n = Message from Remote Hub "n"




Not to scale. Field lengths and interleaving of Data and FEC bits are not precise and intended for illustrative purposes.

FIGURE 10d.




Not to scale. Field lengths and interleaving of Data and FEC bits are not precise and intended for illustrative purposes.

FIGURE 10e.



Avoiding collision on Alert channel using Multiple Chips
HR n = Message from Remote Hub "n"

FIGURE 11.






FIGURE 12a.





FIGURE 12b.



INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2014/016538**A. CLASSIFICATION OF SUBJECT MATTER****H04L 12/26(2006.01)i, H04L 12/24(2006.01)j**

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

H04L 12/26; G05F 5/00; G08C 15/06; G06F 19/00; H04W 56/00; H04W 72/04; H02J 3/00; H04L 12/24

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Korean utility models and applications for utility models
Japanese utility models and applications for utility models

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

eKOMPASS(KIPO internal) & Keywords: substation, improving performance, remote hub

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 2012-0310424 A1 (JEFFREY D. TAFT) 06 December 2012 See abstract, paragraphs [0055]-[0058], [0077]-[0086], [0098], [0104]-[0106], [0120]-[0134] and figure 2.	1-27
Y	JP 2013-005718 A (GENERAL ELECTRIC CO. <GE>) 07 January 2013 See abstract, paragraphs [0013], [0018] and figure 2.	1-14
A		15-27
Y	US 2010-0286840 A1 (PHILLIP W. POWELL et al.) 11 November 2010 See abstract, paragraphs [0075]-[0089], [0098], [0122] and figures 1-2.	11-13, 15-27
A	US 2010-0188260 A1 (MARK K. CORNWALL et al.) 29 July 2010 See paragraphs [0033]-[0042] and figures 1, 5.	1-27
A	US 2013-0034086 A1 (MARCIA R. MARTIN et al.) 07 February 2013 See paragraphs [0026]-[0031], [0056]-[0061] and figure 2.	1-27

 Further documents are listed in the continuation of Box C. See patent family annex.

* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family


Date of the actual completion of the international search

13 June 2014 (13.06.2014)

Date of mailing of the international search report

16 June 2014 (16.06.2014)

Name and mailing address of the ISA/KR


 International Application Division
Korean Intellectual Property Office
189 Cheongsa-ro, Seo-gu, Daejeon Metropolitan City, 302-701,
Republic of Korea

Facsimile No. +82-42-472-7140

Authorized officer

KIM, Seong Woo

Telephone No. +82-42-481-3348



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2014/016538

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2012-0310424 A1	06/12/2012	US 2012-310434 A1 US 2012-310435 A1 US 2012-310558 A1	06/12/2012 06/12/2012 06/12/2012
JP 2013-005718 A	07/01/2013	AU 2012-203447 A1 BR 102012014647 A2 CA 2779972 A1 DE 102012105184 A1 GB 2491952 A NZ 600601 A US 2012-0323388 A1	10/01/2013 17/12/2013 16/12/2012 20/12/2012 19/12/2012 20/12/2013 20/12/2012
US 2010-0286840 A1	11/11/2010	AU 2010-245937 A1 CA 2758536 A1 CN 102549880 A EP 2427949 A2 EP 2427949 A4 IL 216179 D0 JP 2012-526515 A JP 2013-243925 A JP 5312686 B2 KR 10-1353005 B1 KR 10-2012-0023620 A MX 2011009052 A RU 2480885 C1 US 2014-0008982 A1 US 8437883 B2 US 8577510 B2 WO 2010-129691 A2 WO 2010-129691 A3	01/09/2011 11/11/2010 04/07/2012 14/03/2012 14/08/2013 31/01/2012 25/10/2012 05/12/2013 09/10/2013 21/01/2014 13/03/2012 28/02/2012 27/04/2013 09/01/2014 07/05/2013 05/11/2013 11/11/2010 10/03/2011
US 2010-0188260 A1	29/07/2010	CA 2688741 A1 CA 2689772 A1 CA 2689773 A1 US 2010-0188263 A1 US 2010-0265095 A1 US 8242887 B2 US 8310341 B2 US 8436744 B2	29/07/2010 29/07/2010 29/07/2010 29/07/2010 21/10/2010 14/08/2012 13/11/2012 07/05/2013
US 2013-0034086 A1	07/02/2013	WO 2013-020053 A1	07/02/2013